

US 20250046179A1

(19) **United States**

(12) **Patent Application Publication**  
**PAULRAJ et al.**

(10) **Pub. No.: US 2025/0046179 A1**

(43) **Pub. Date: Feb. 6, 2025**

(54) **SYSTEM AND METHOD FOR TELEMETRY STREAMING AS A SERVICE**

(71) Applicant: **Dell Products L.P.**, Round Rock, TX (US)

(72) Inventors: **DEEPAGANESH PAULRAJ**, Bangalore (IN); **RAVI SHEKHAR SINGH**, Bangalore (IN); **VAISHNAVI SUCHINDRAN**, Bangalore (IN); **SHABU VELLARAKKATTU CHALIL**, Bangalore (IN); **SAYALI SUNIL NIKAM**, Maharashtra (IN); **HIREN KISHORBHAI PITRODA**, Gujarat (IN)

(21) Appl. No.: **18/477,164**

(22) Filed: **Sep. 28, 2023**

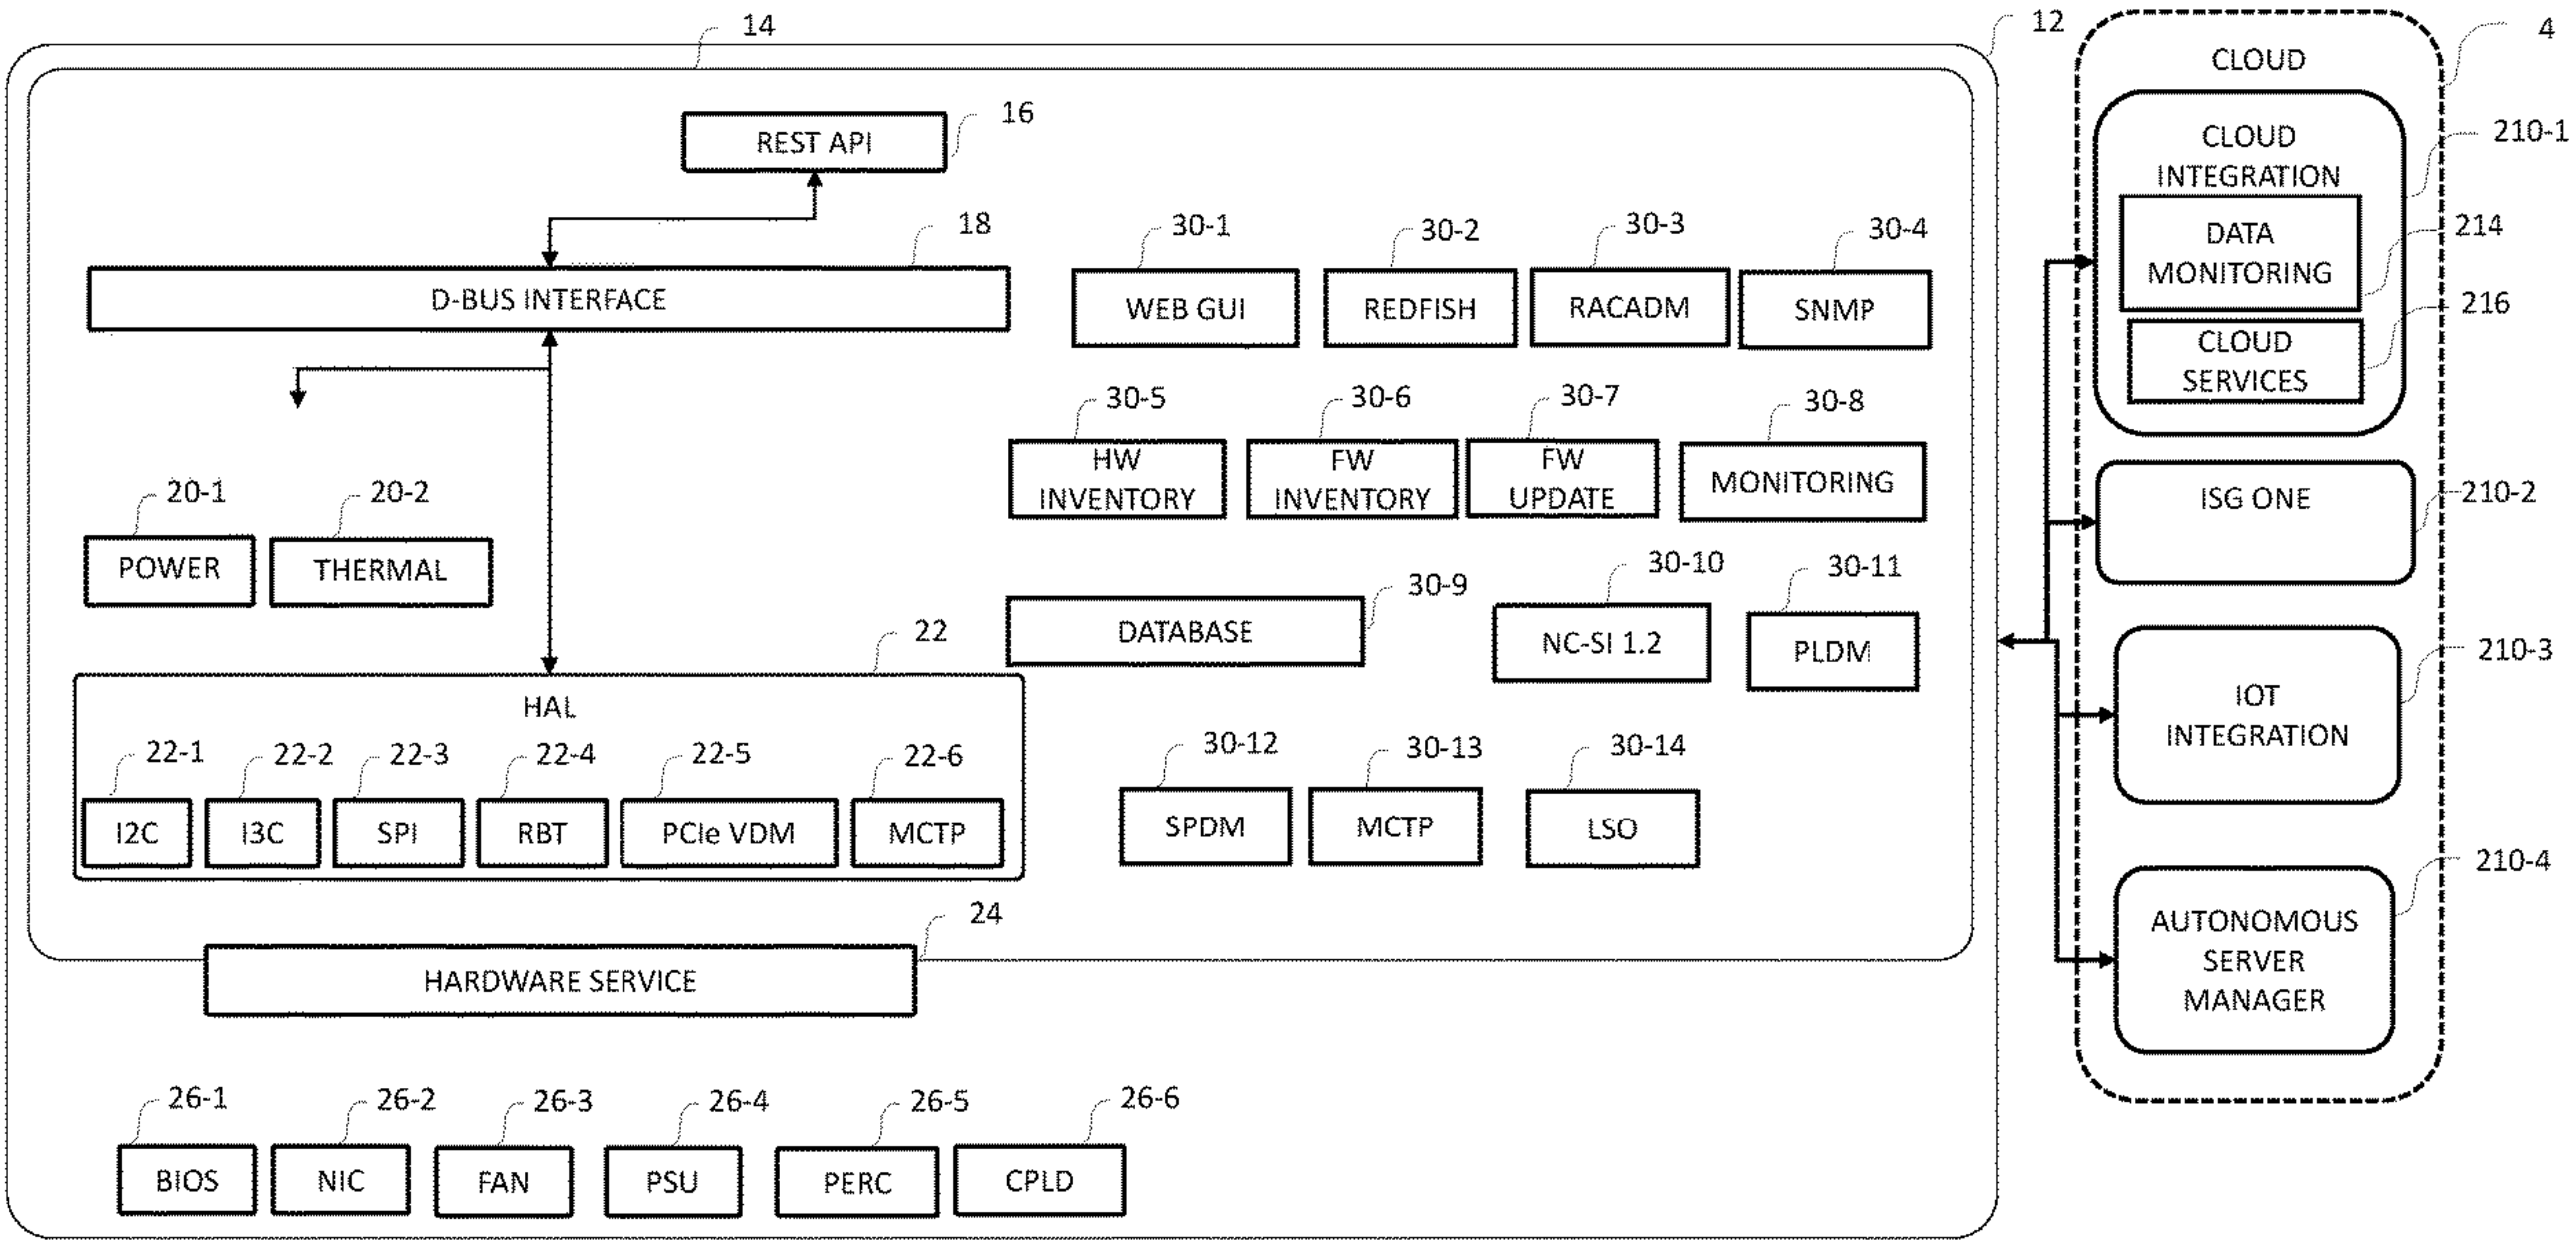
(30) **Foreign Application Priority Data**  
Aug. 4, 2023 (IN) ..... 202311052573

**Publication Classification**

(51) **Int. Cl.**  
**G08C 19/00** (2006.01)  
**H04L 67/12** (2006.01)

(52) **U.S. Cl.**  
CPC ..... **G08C 19/00** (2013.01); **H04L 67/12** (2013.01)

(57) **ABSTRACT**  
In one or more embodiments, one or more information handling systems may comprise a Remote Access Controller (RAC) configured with a minimum number of built-in features and a private cloud server (PCS) communicatively coupled to the RAC and storing a set of telemetry services. The PCS can execute the telemetry services to monitor and manage telemetry operations of a plurality of servers. The PCS may receive a first set of telemetry data from one or more servers, receive outgoing telemetry traffic associated with the first set of telemetry data, determine one or more telemetry collection requirements and communicate a second set of telemetry data and communicate a second set of telemetry data to the external cloud server, wherein the second set of telemetry data communicated to the external cloud server contains less data than the first set of telemetry data and meets the one or more telemetry collection requirements.



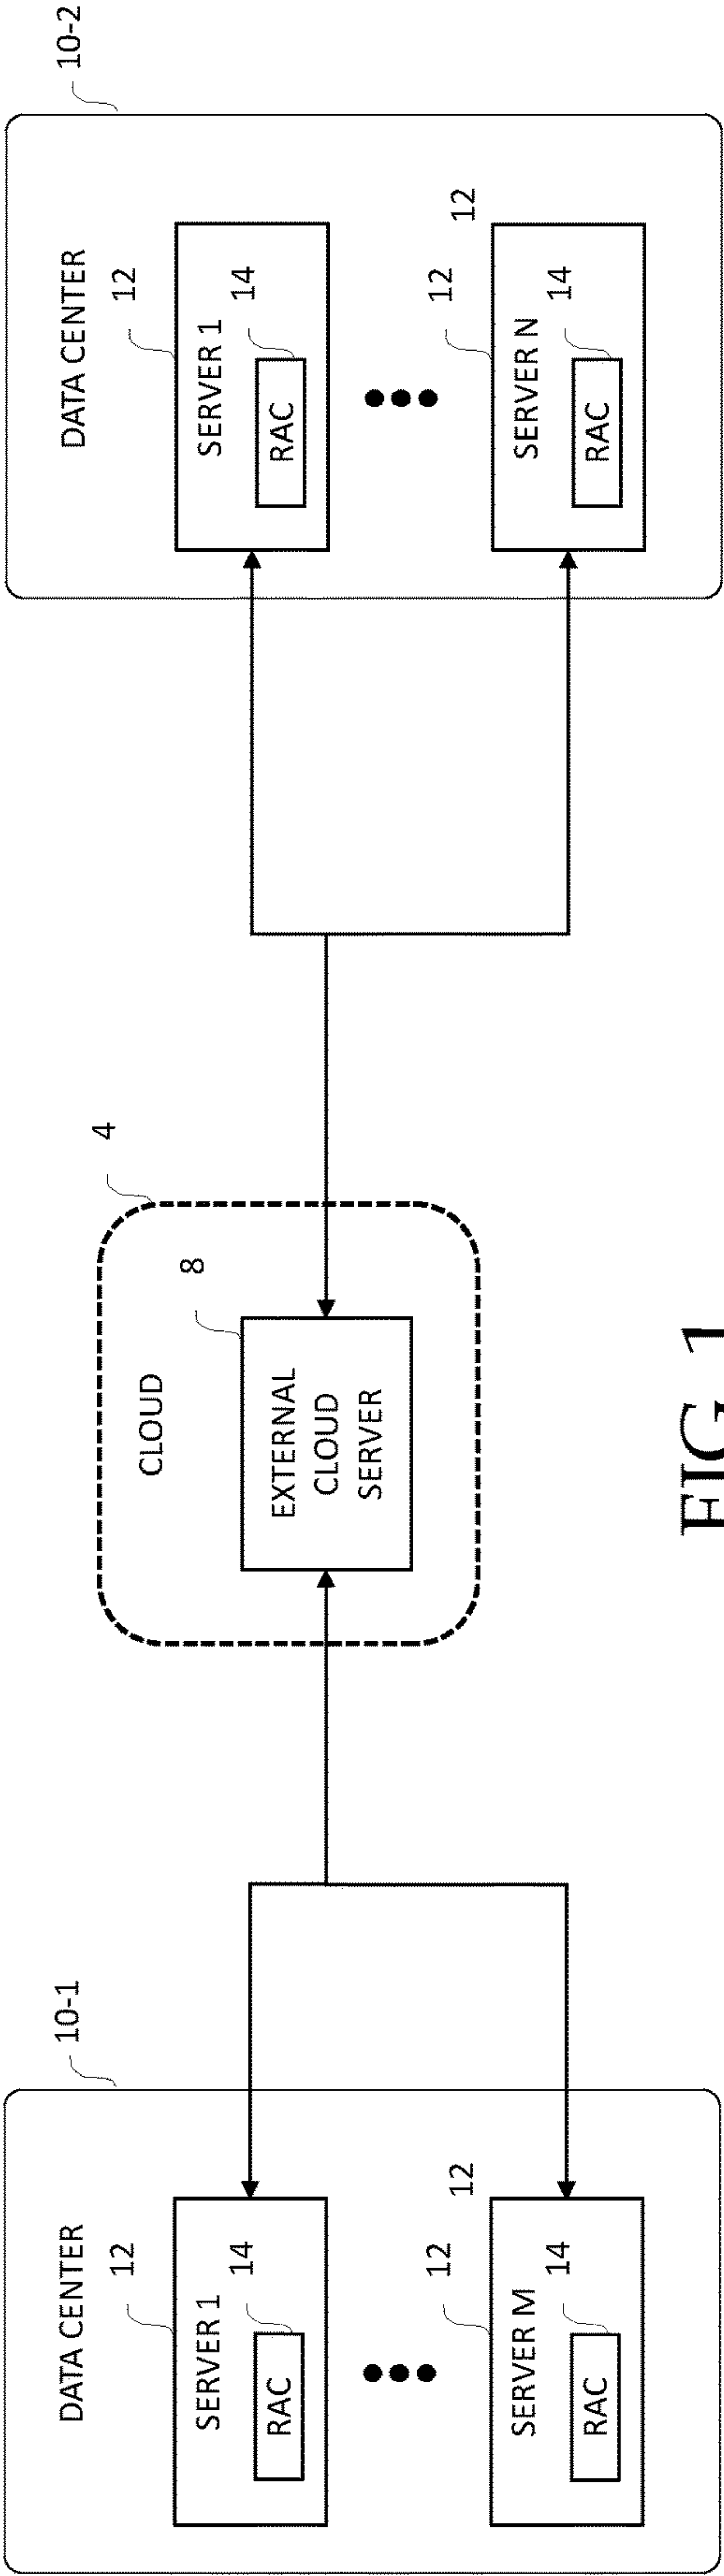


FIG. 1

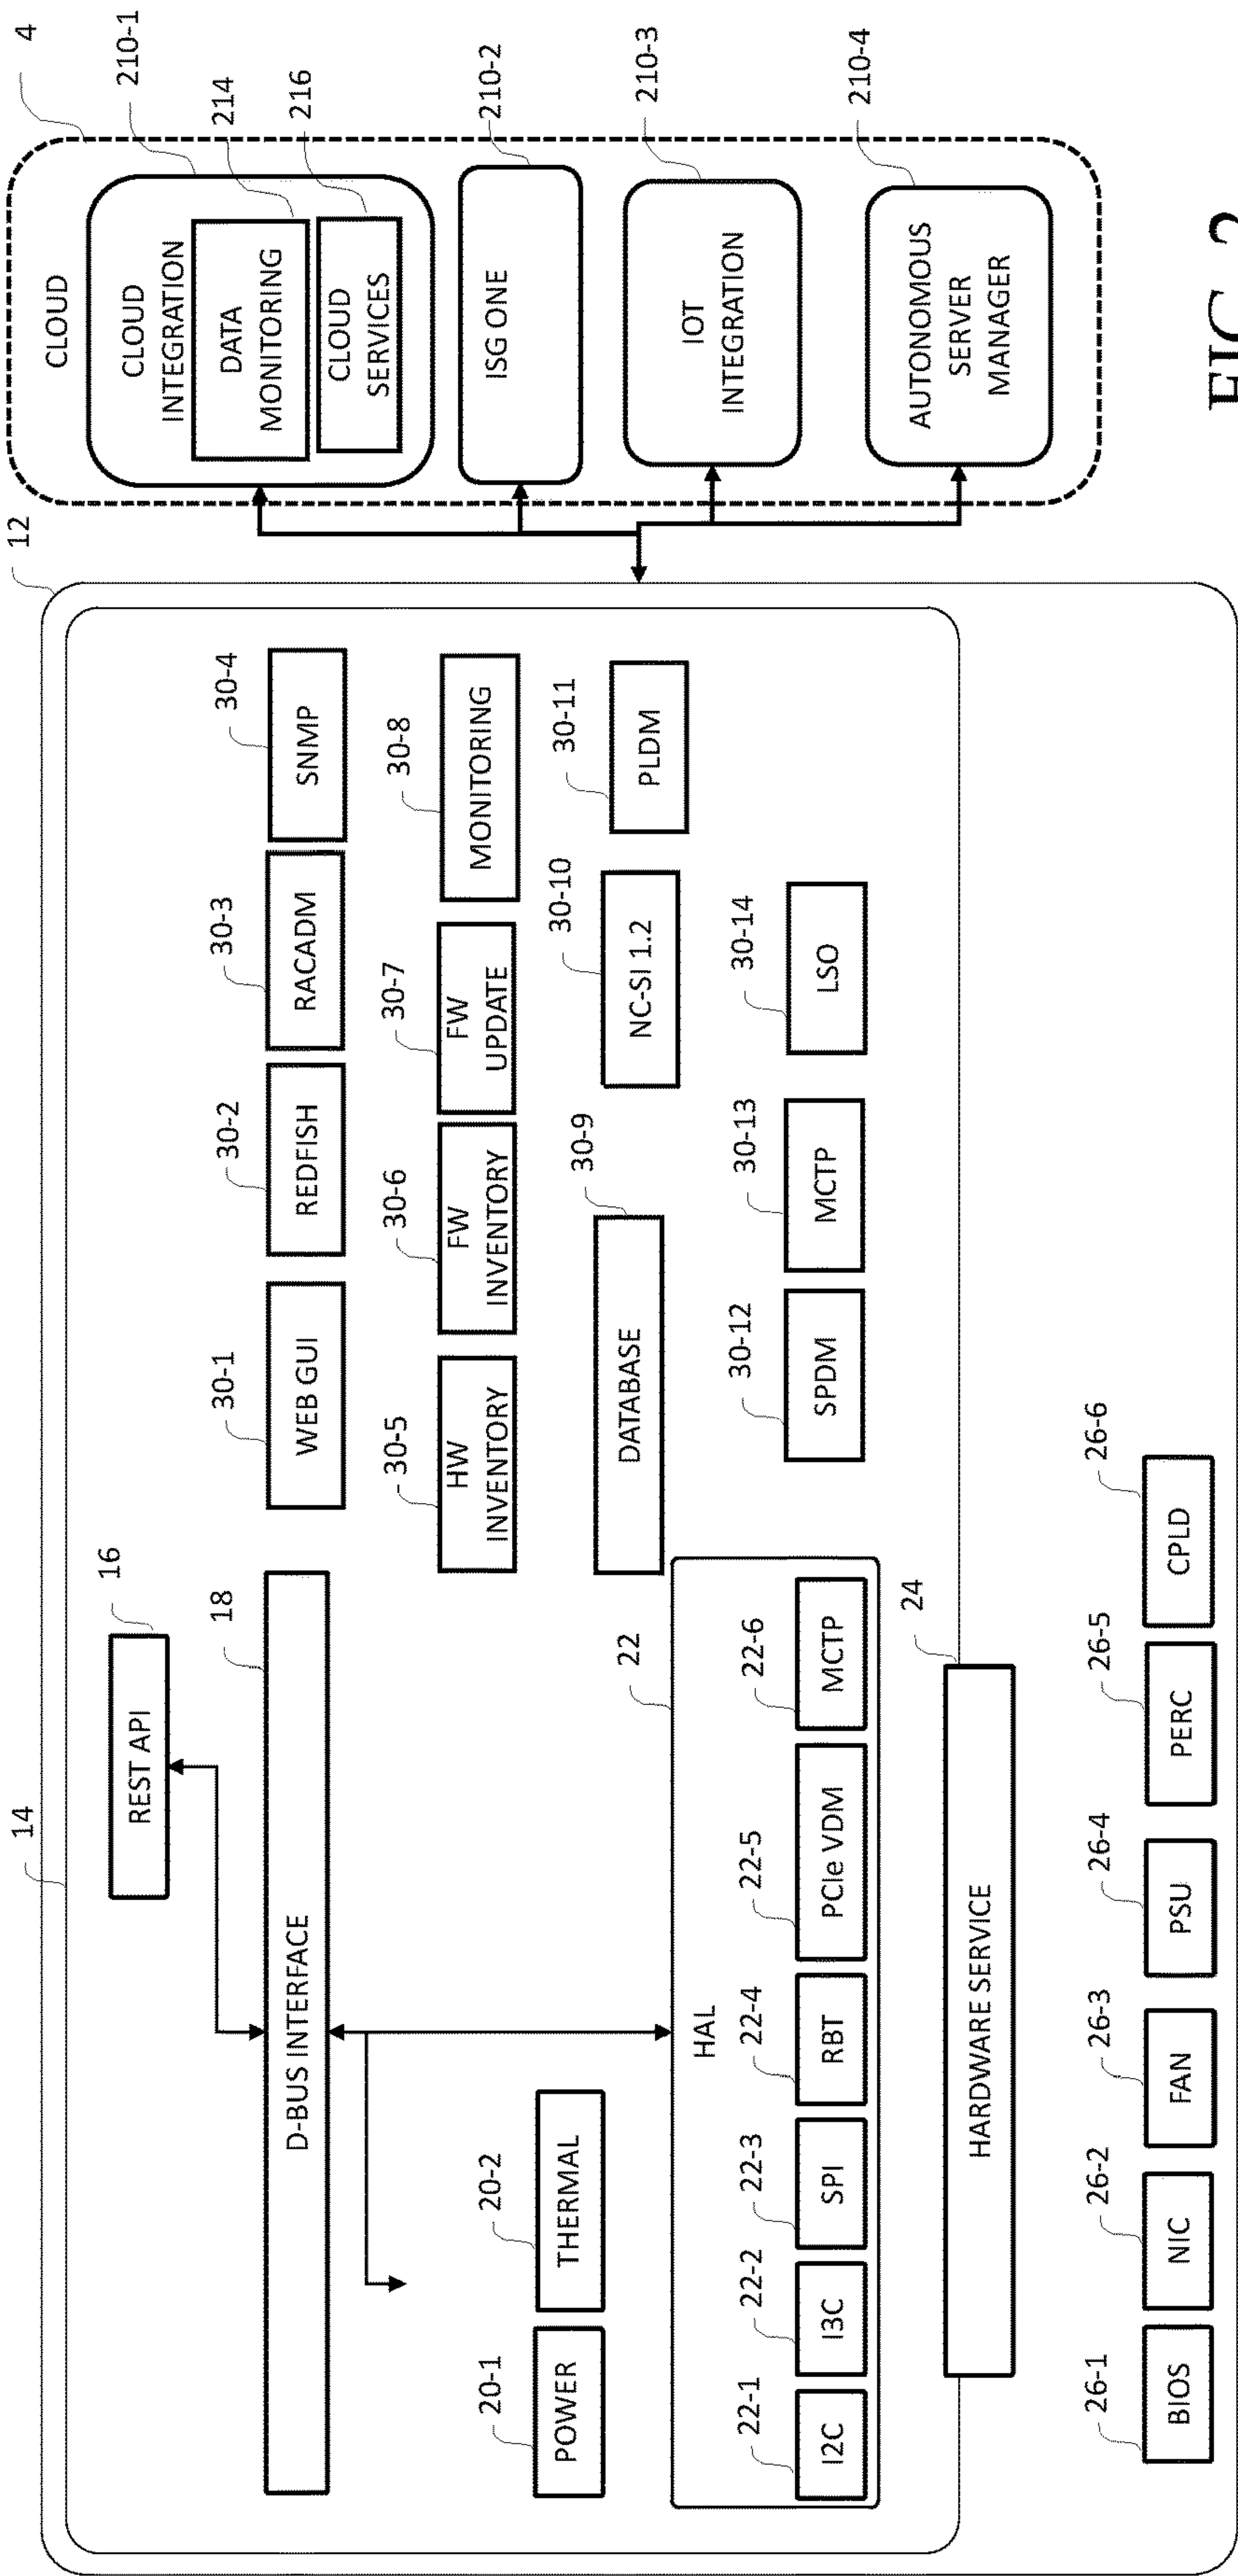


FIG. 2

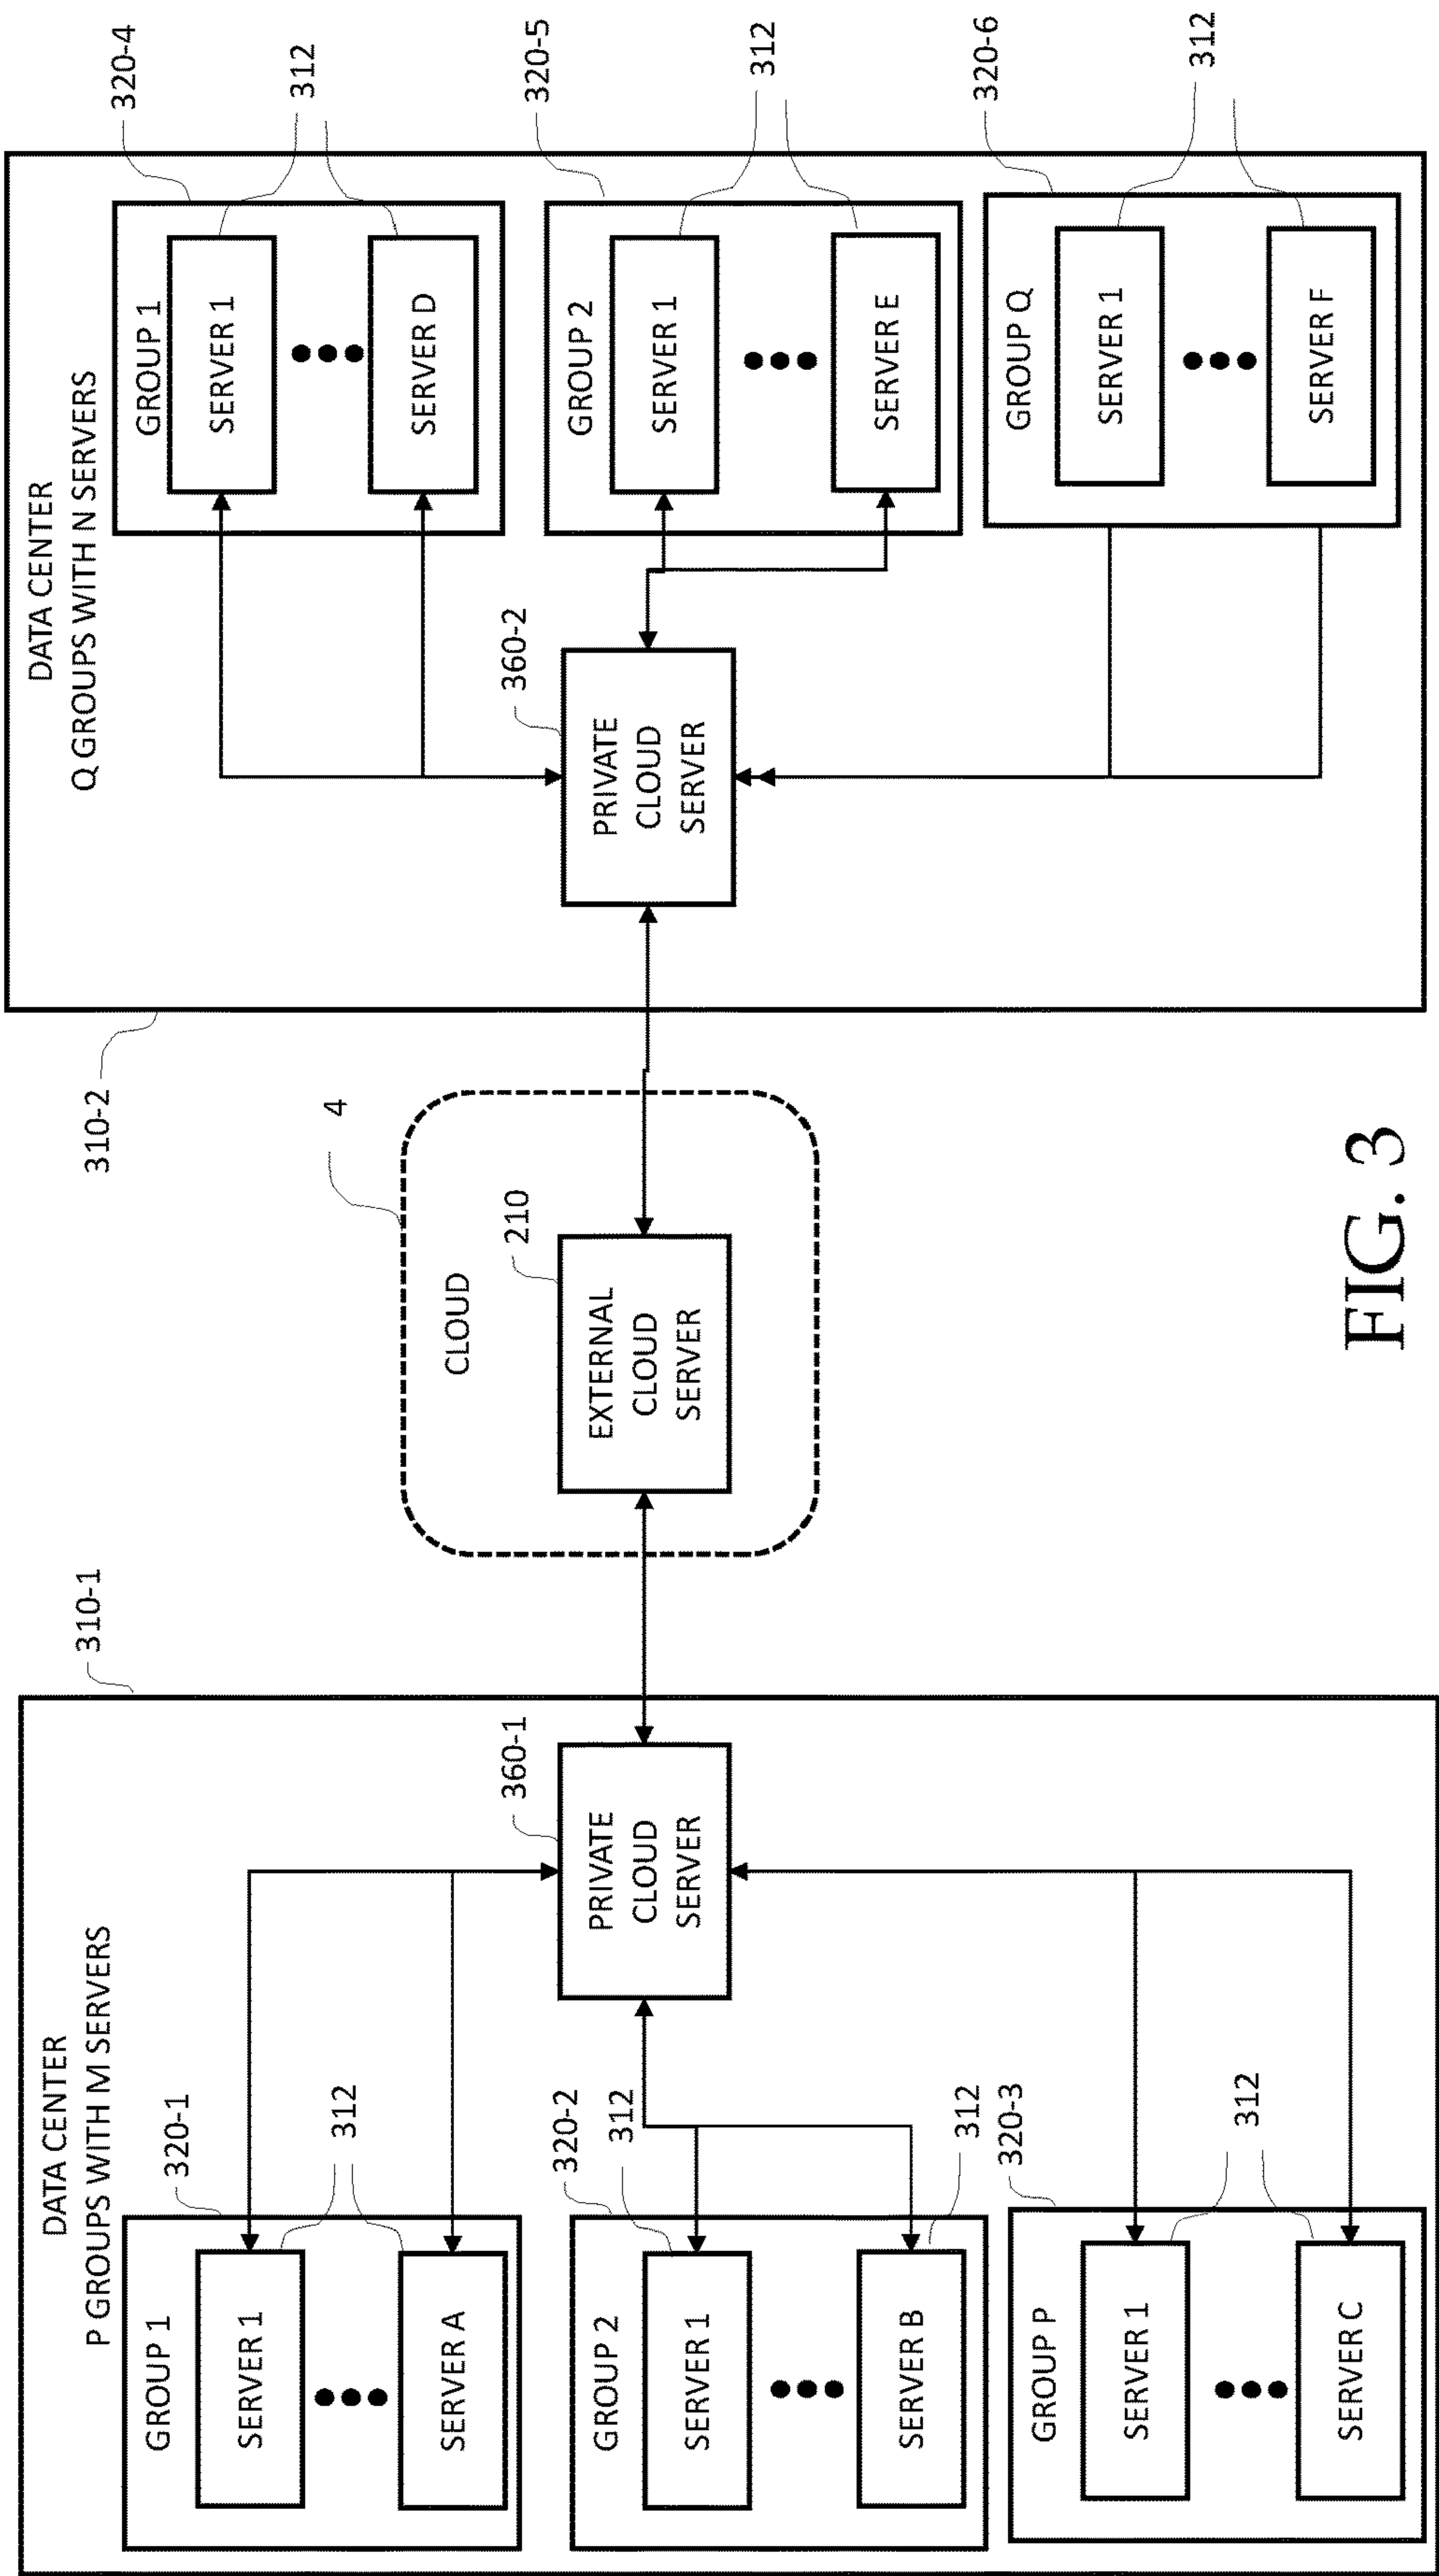
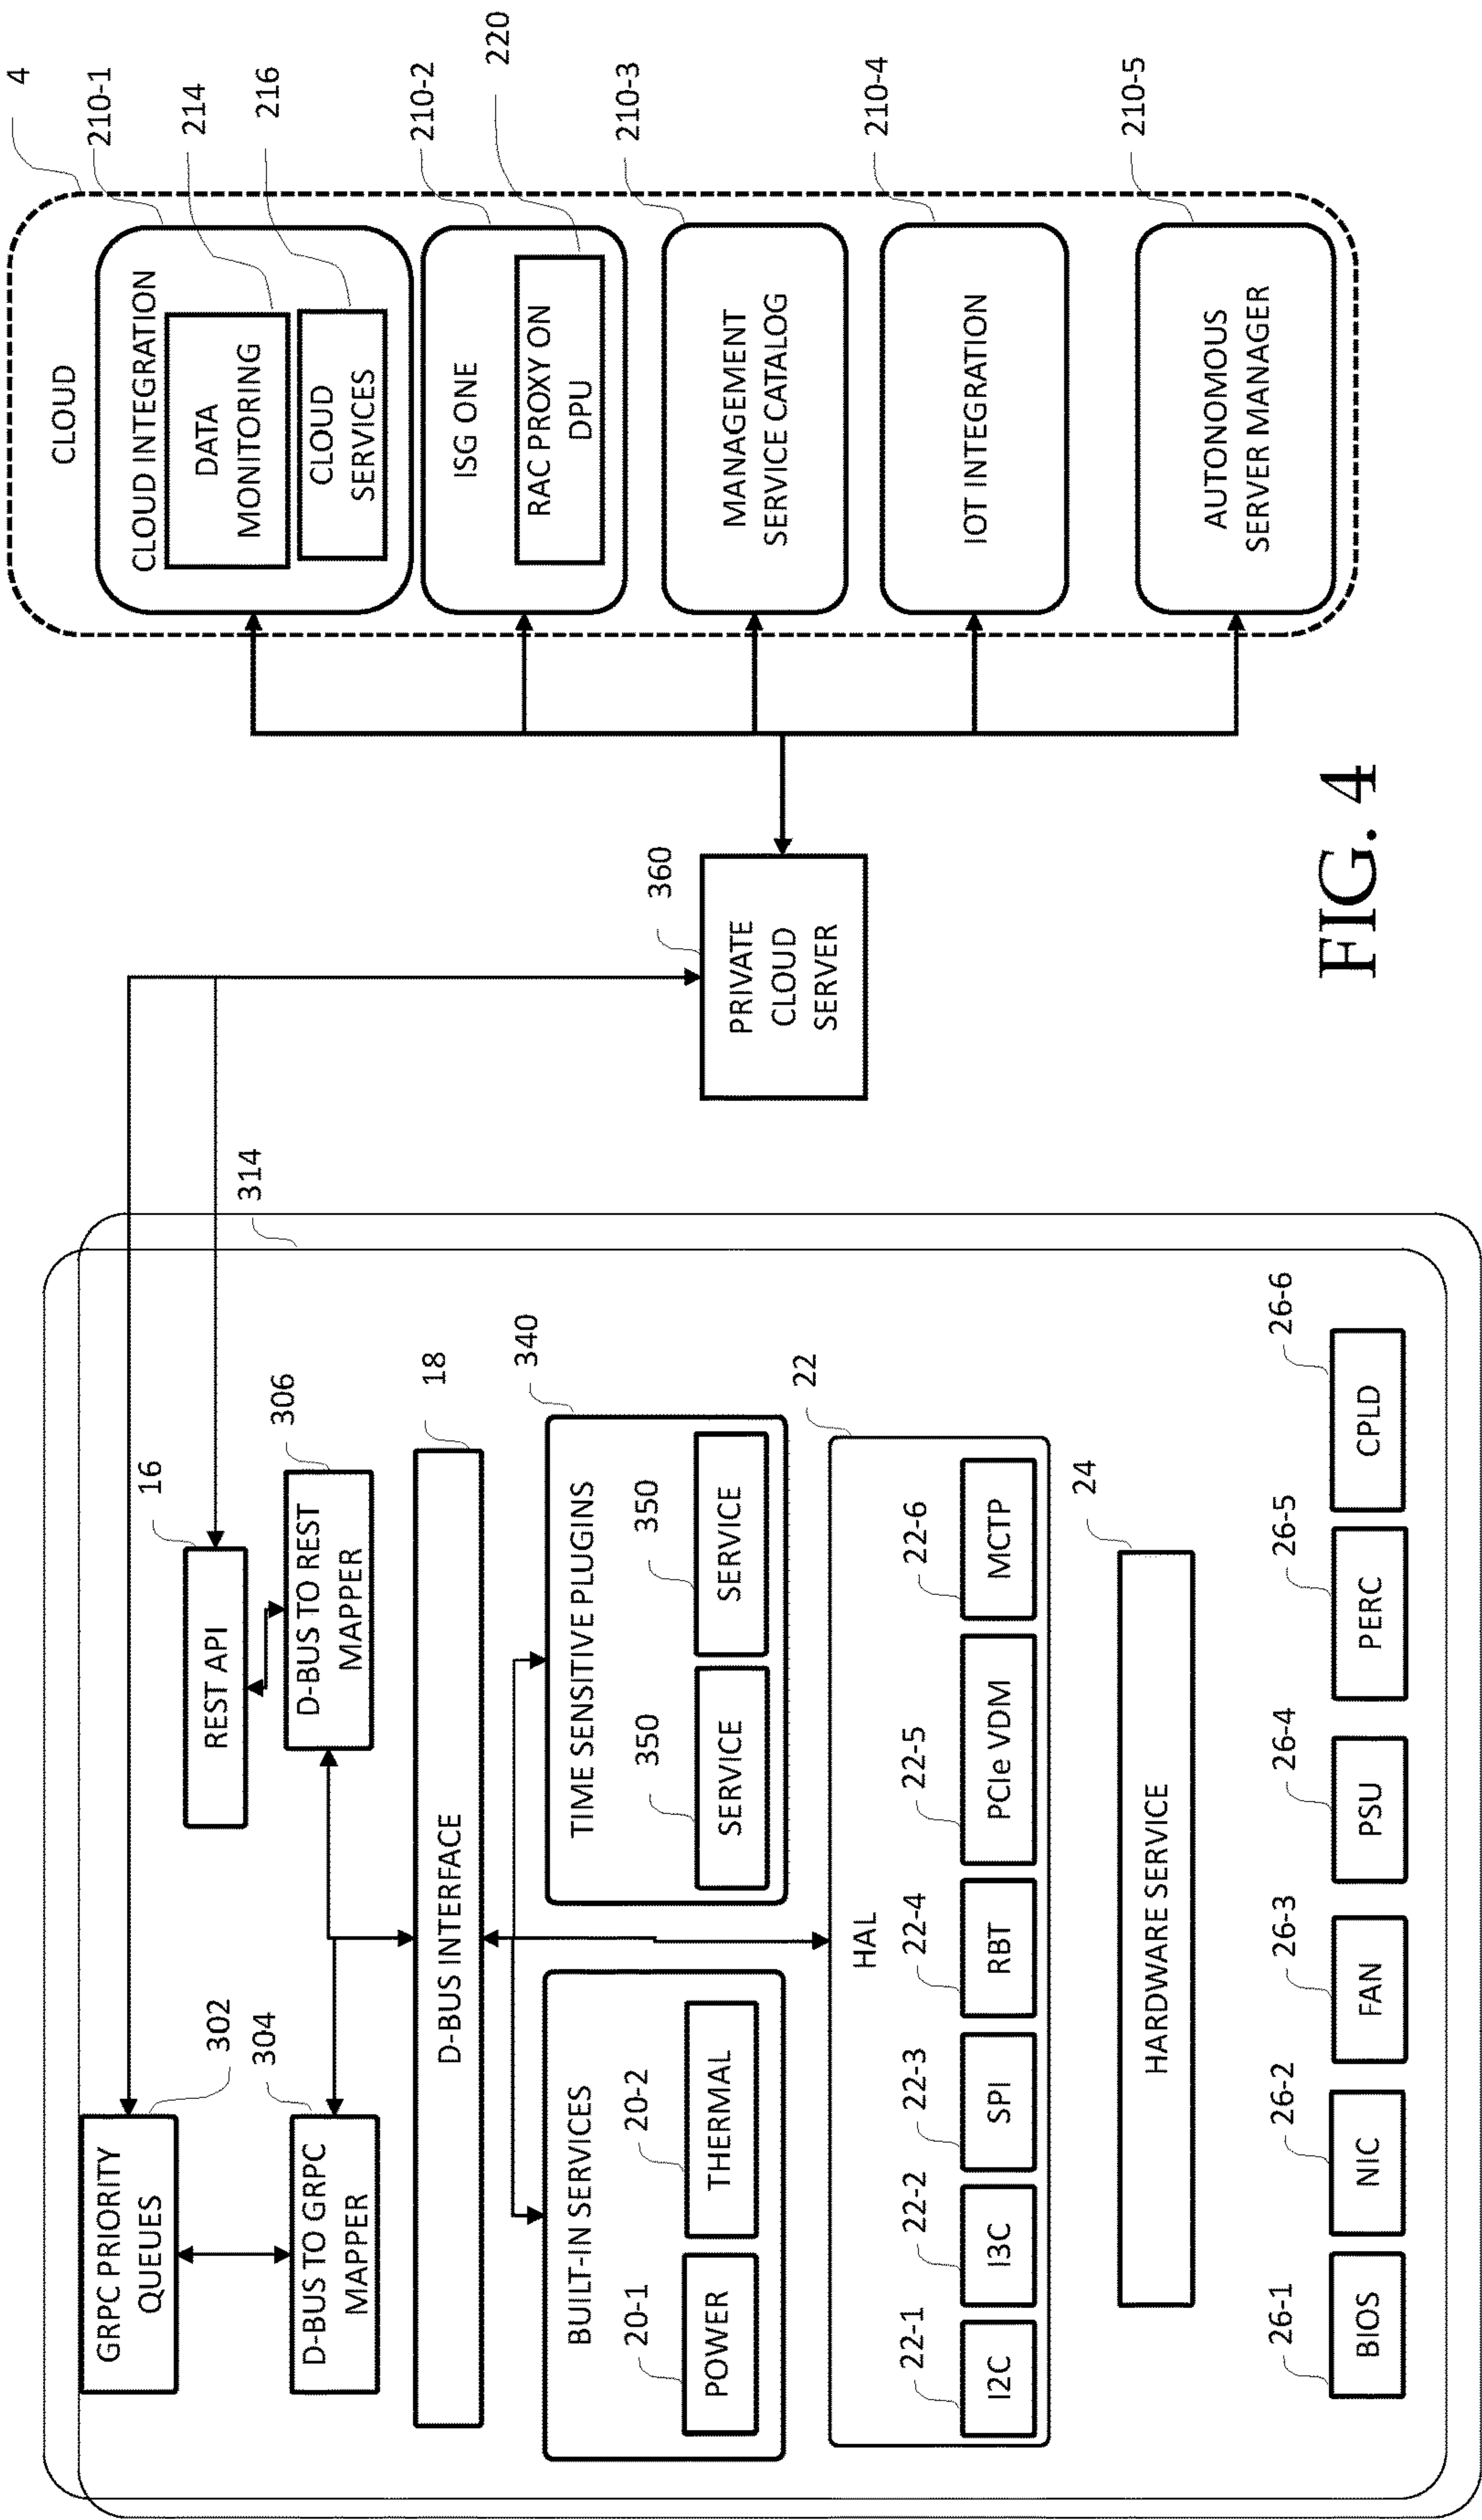


FIG. 3





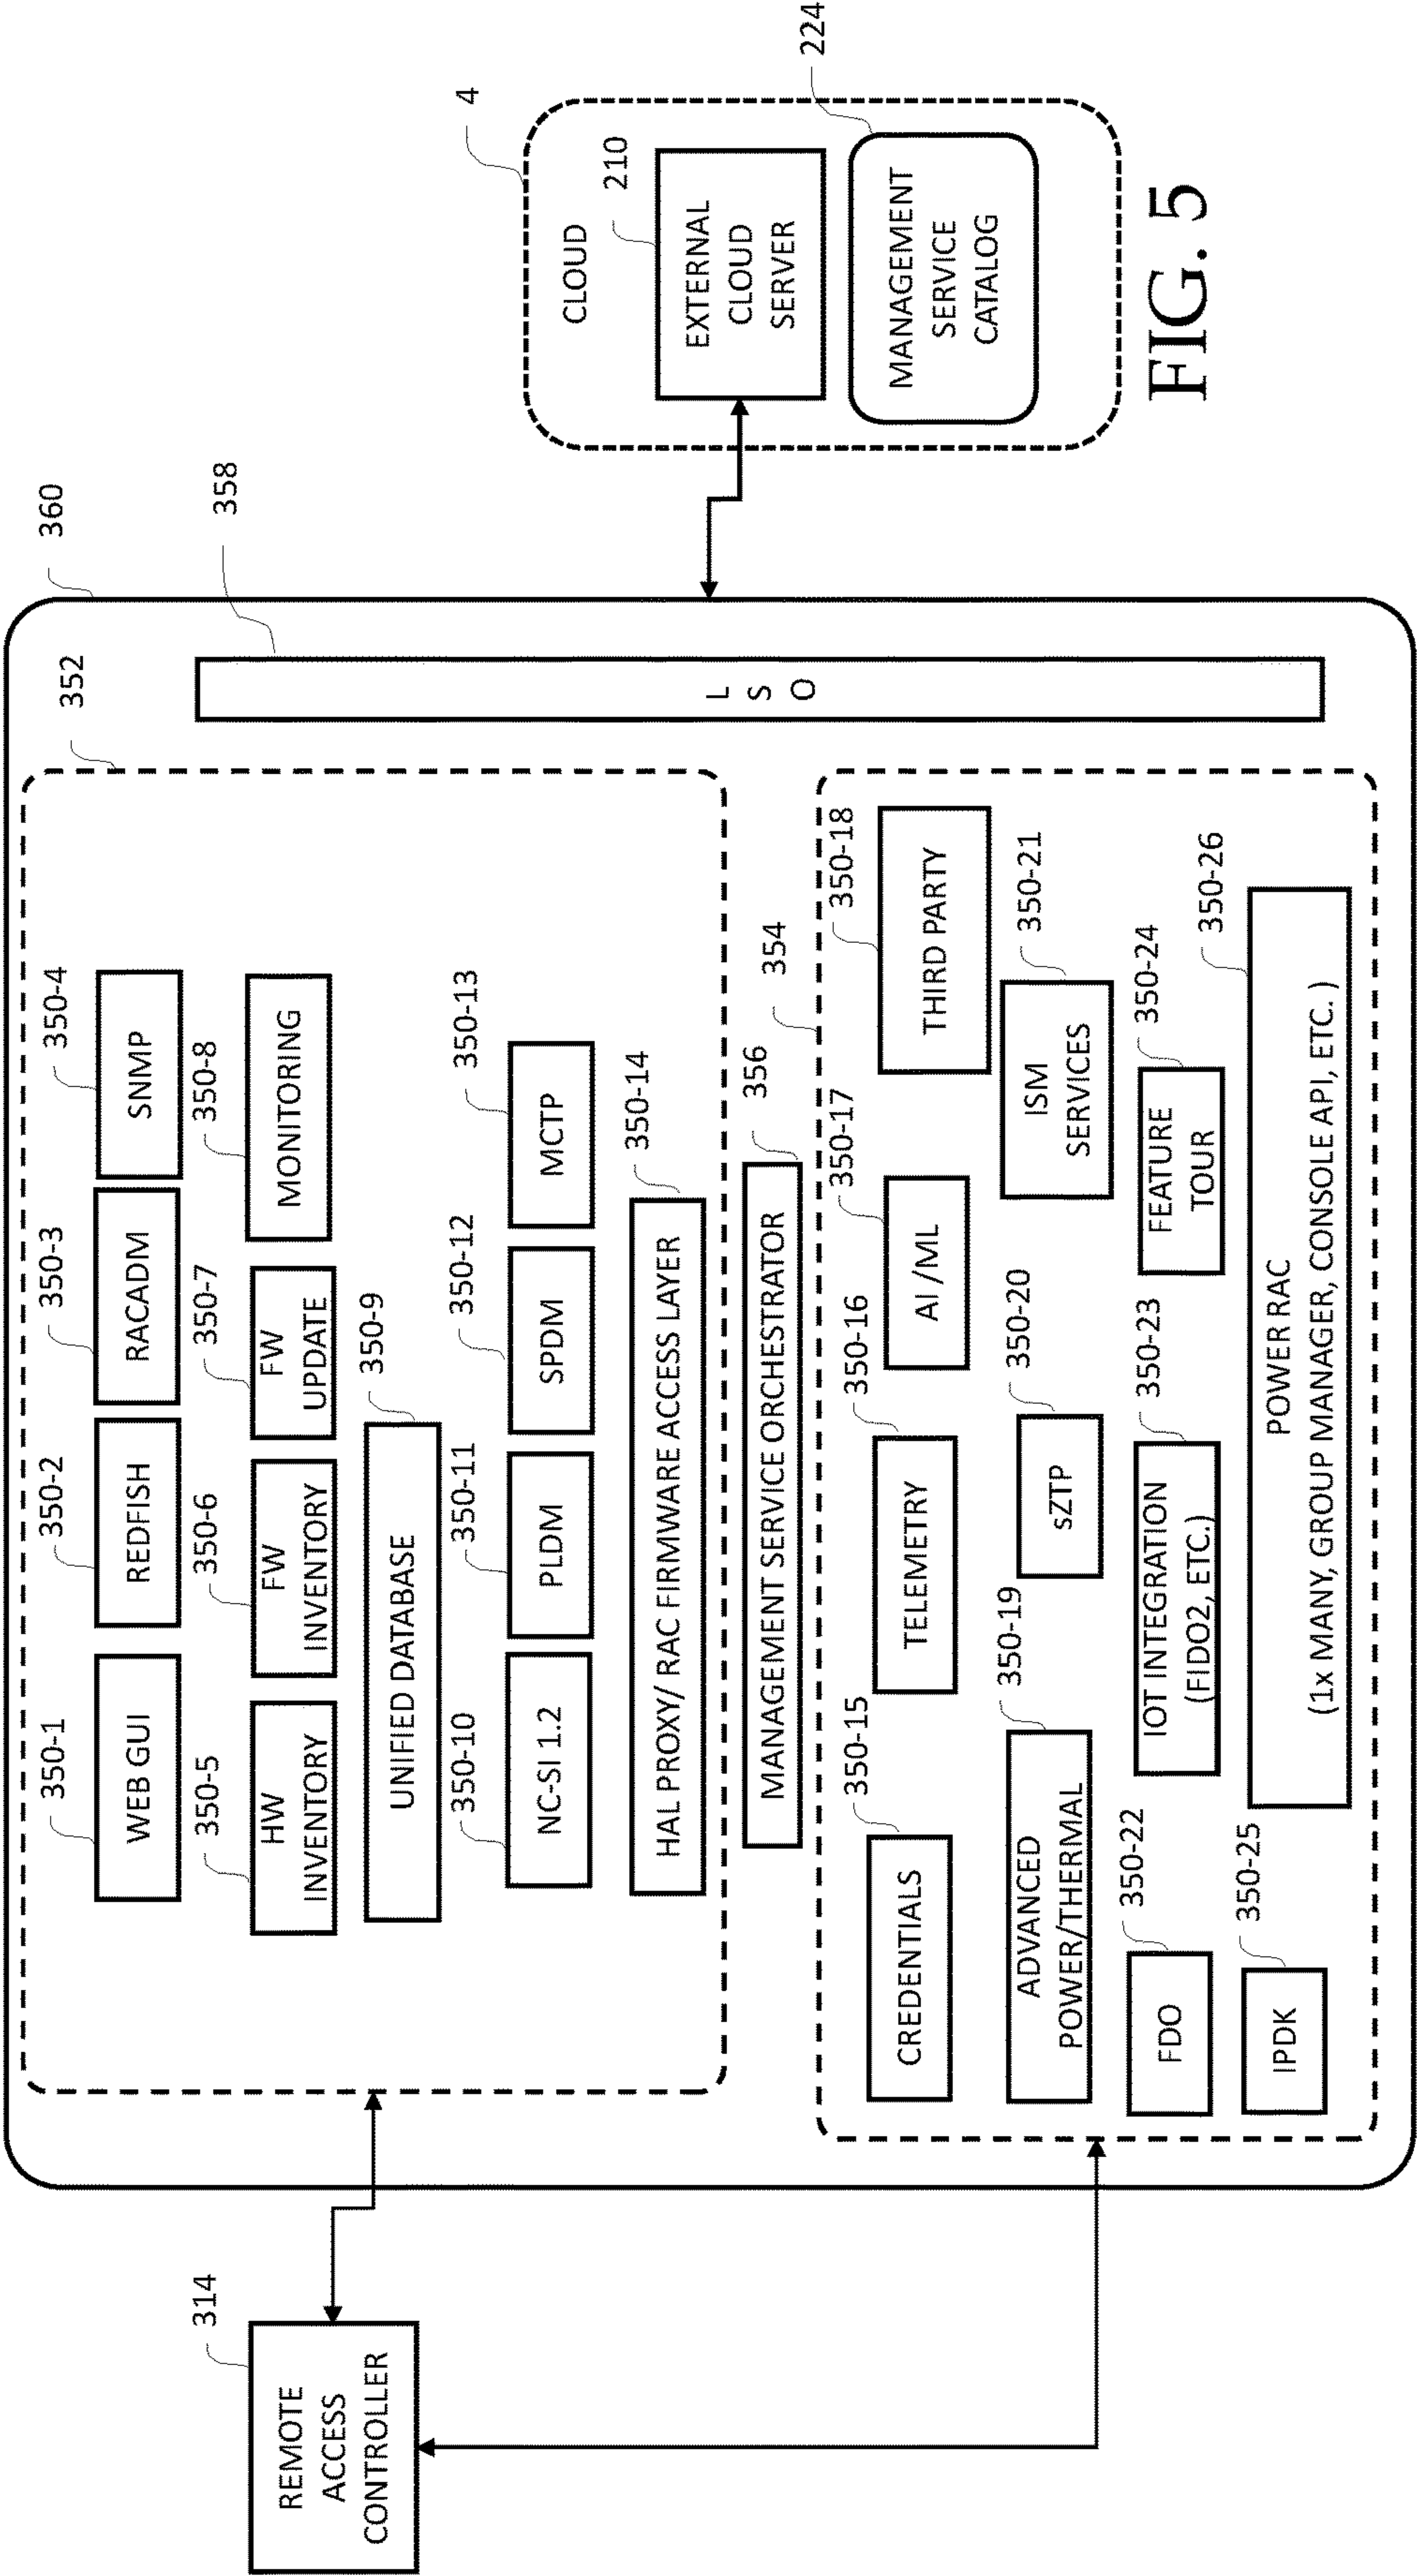


FIG. 5



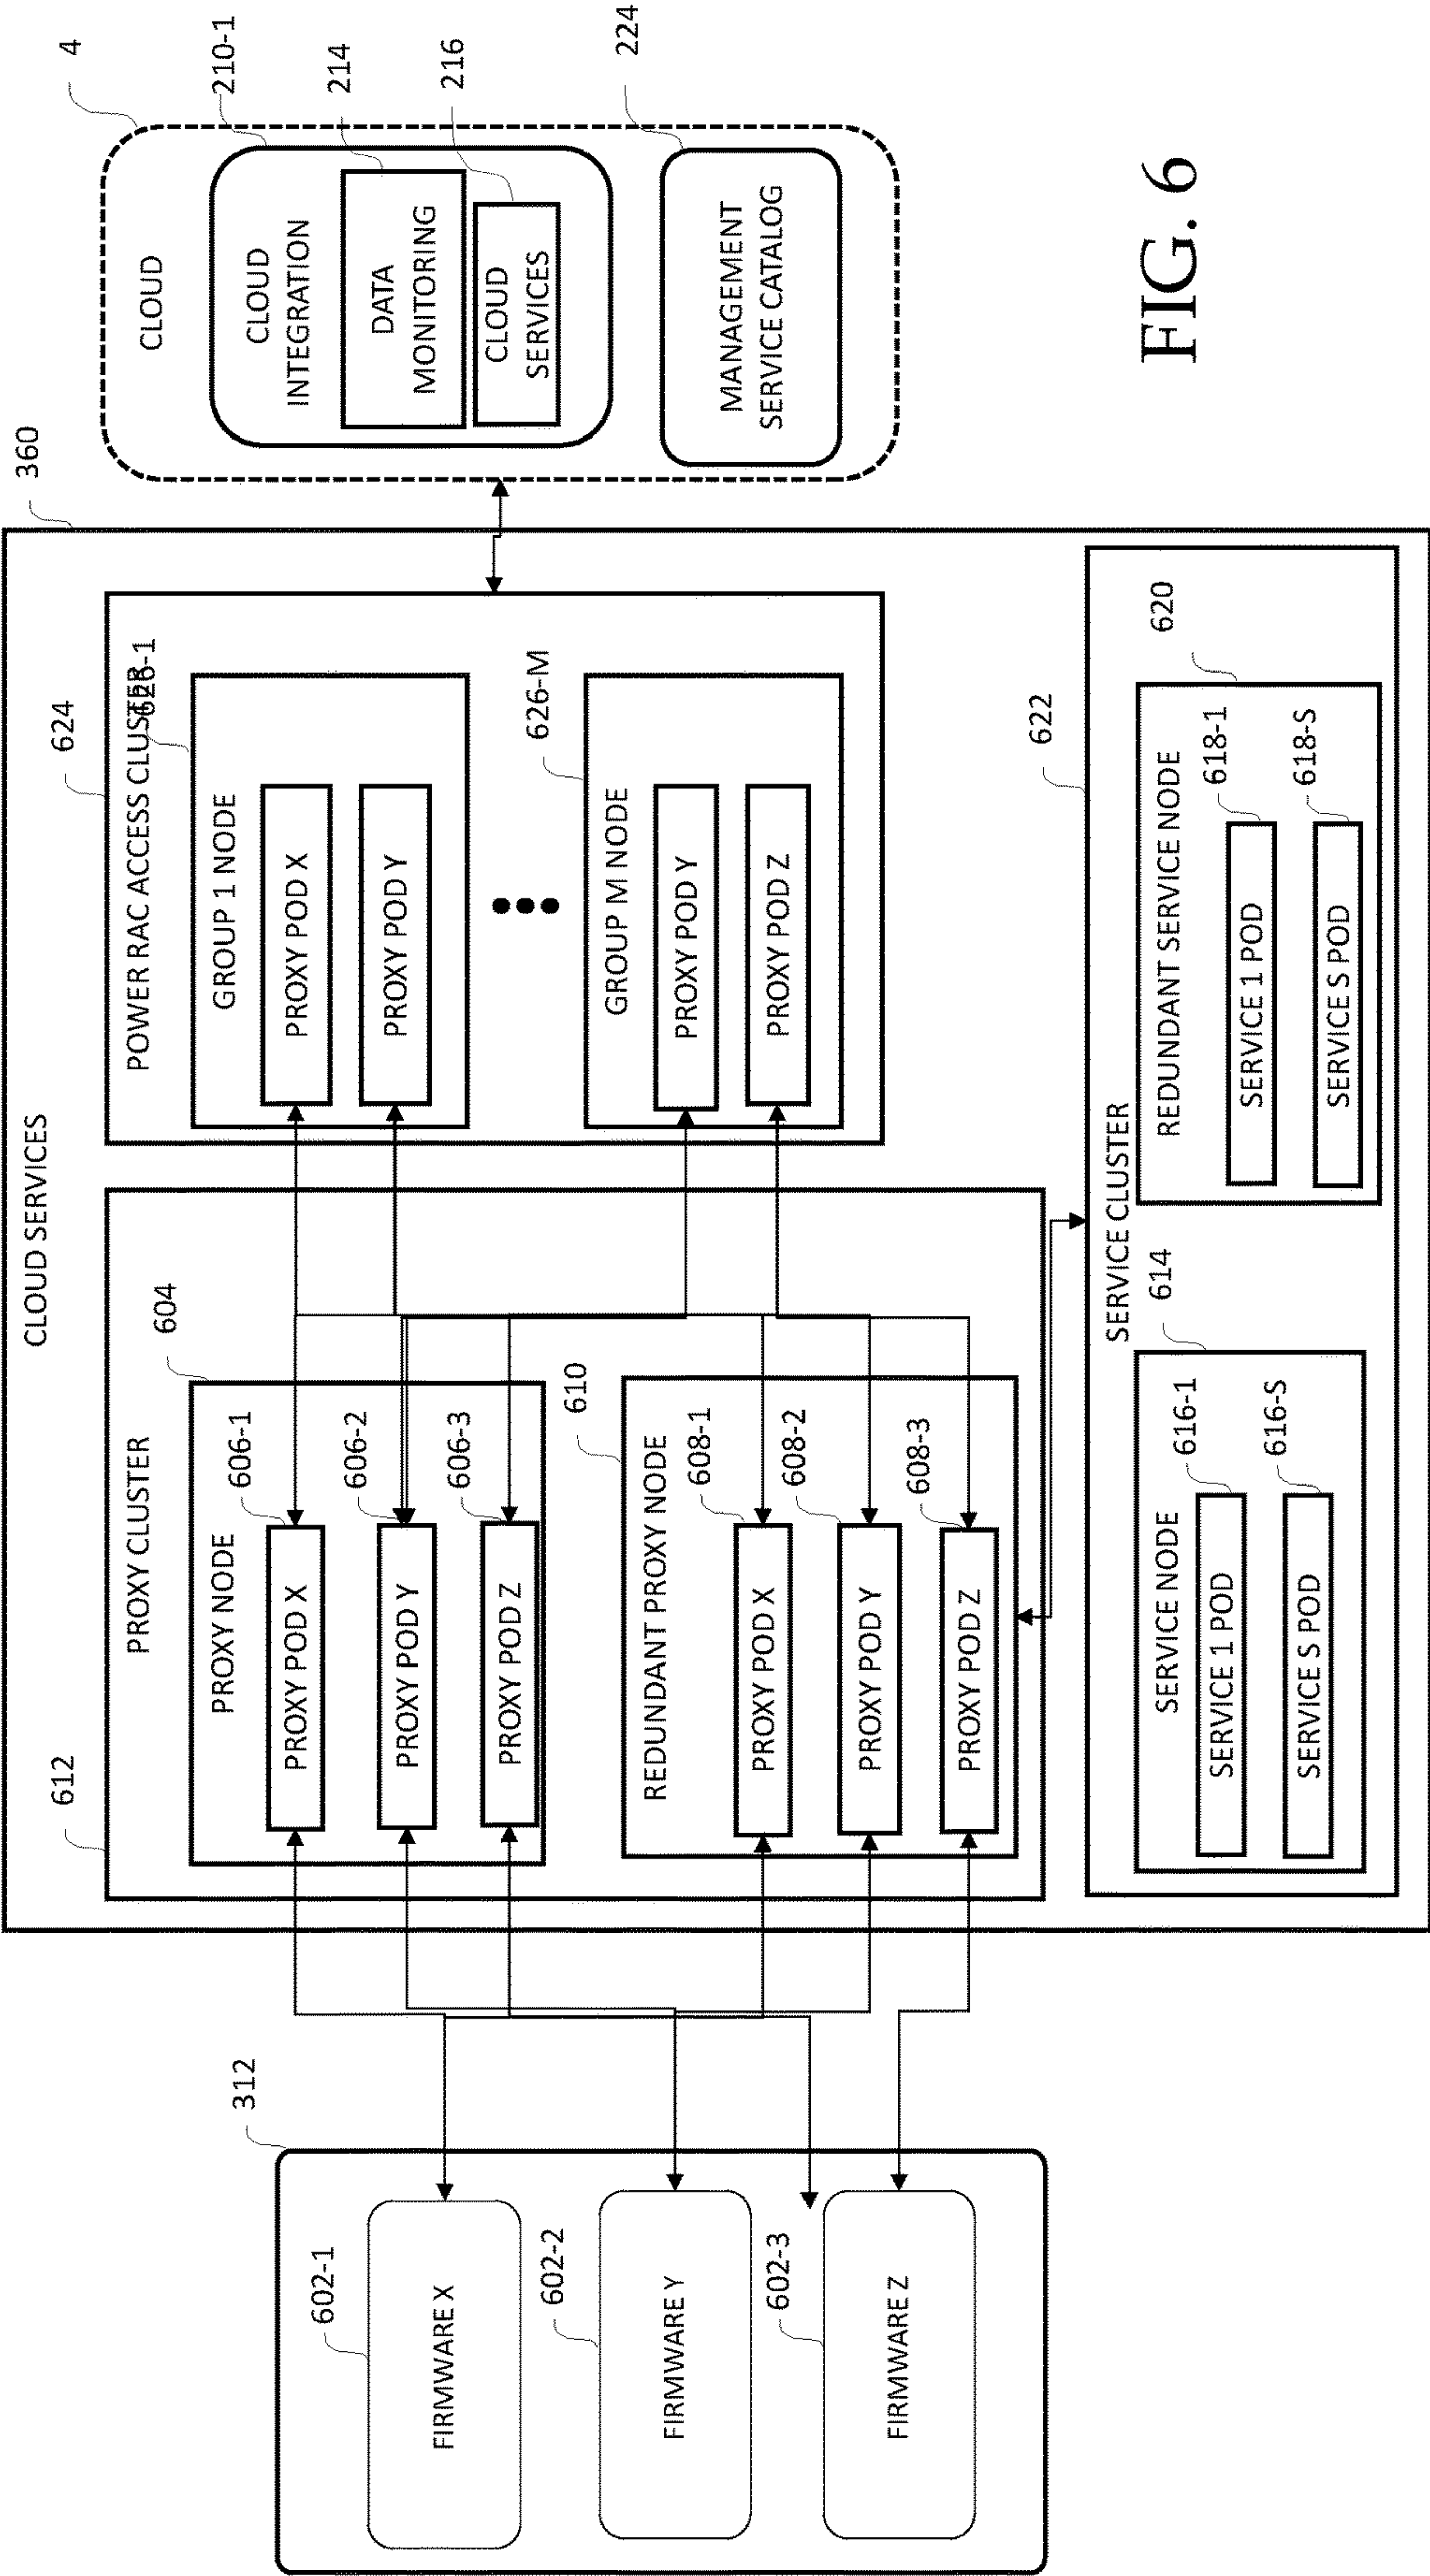
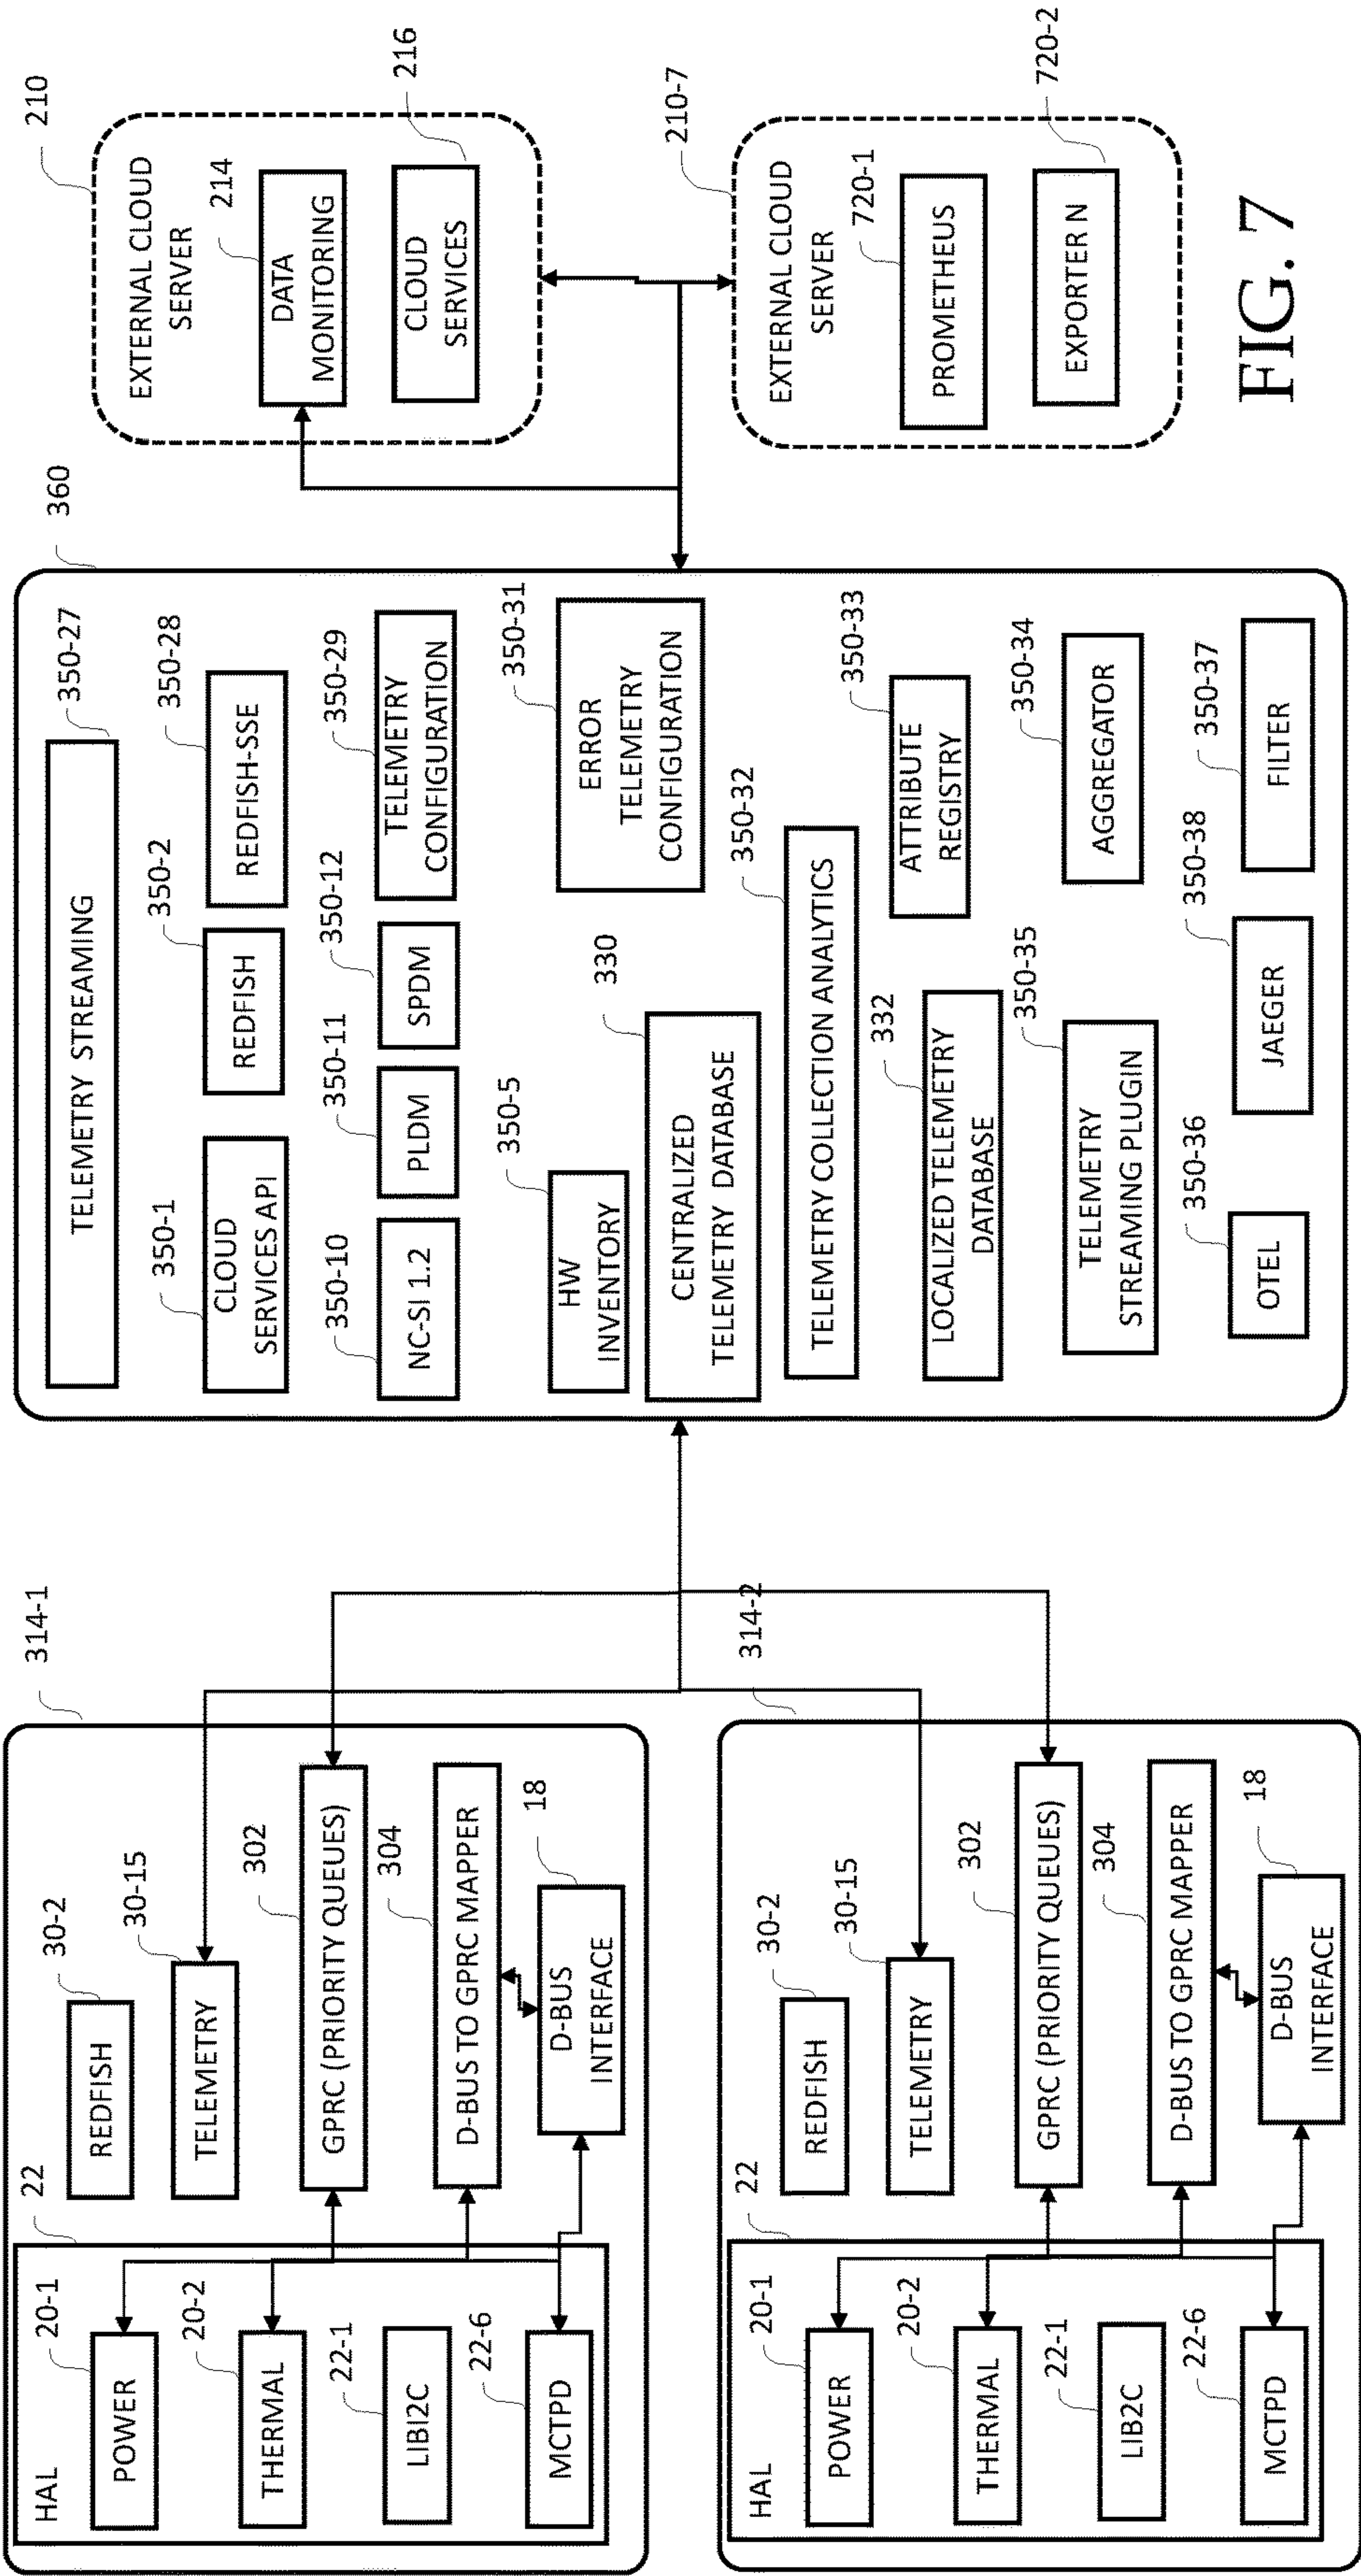


FIG. 6





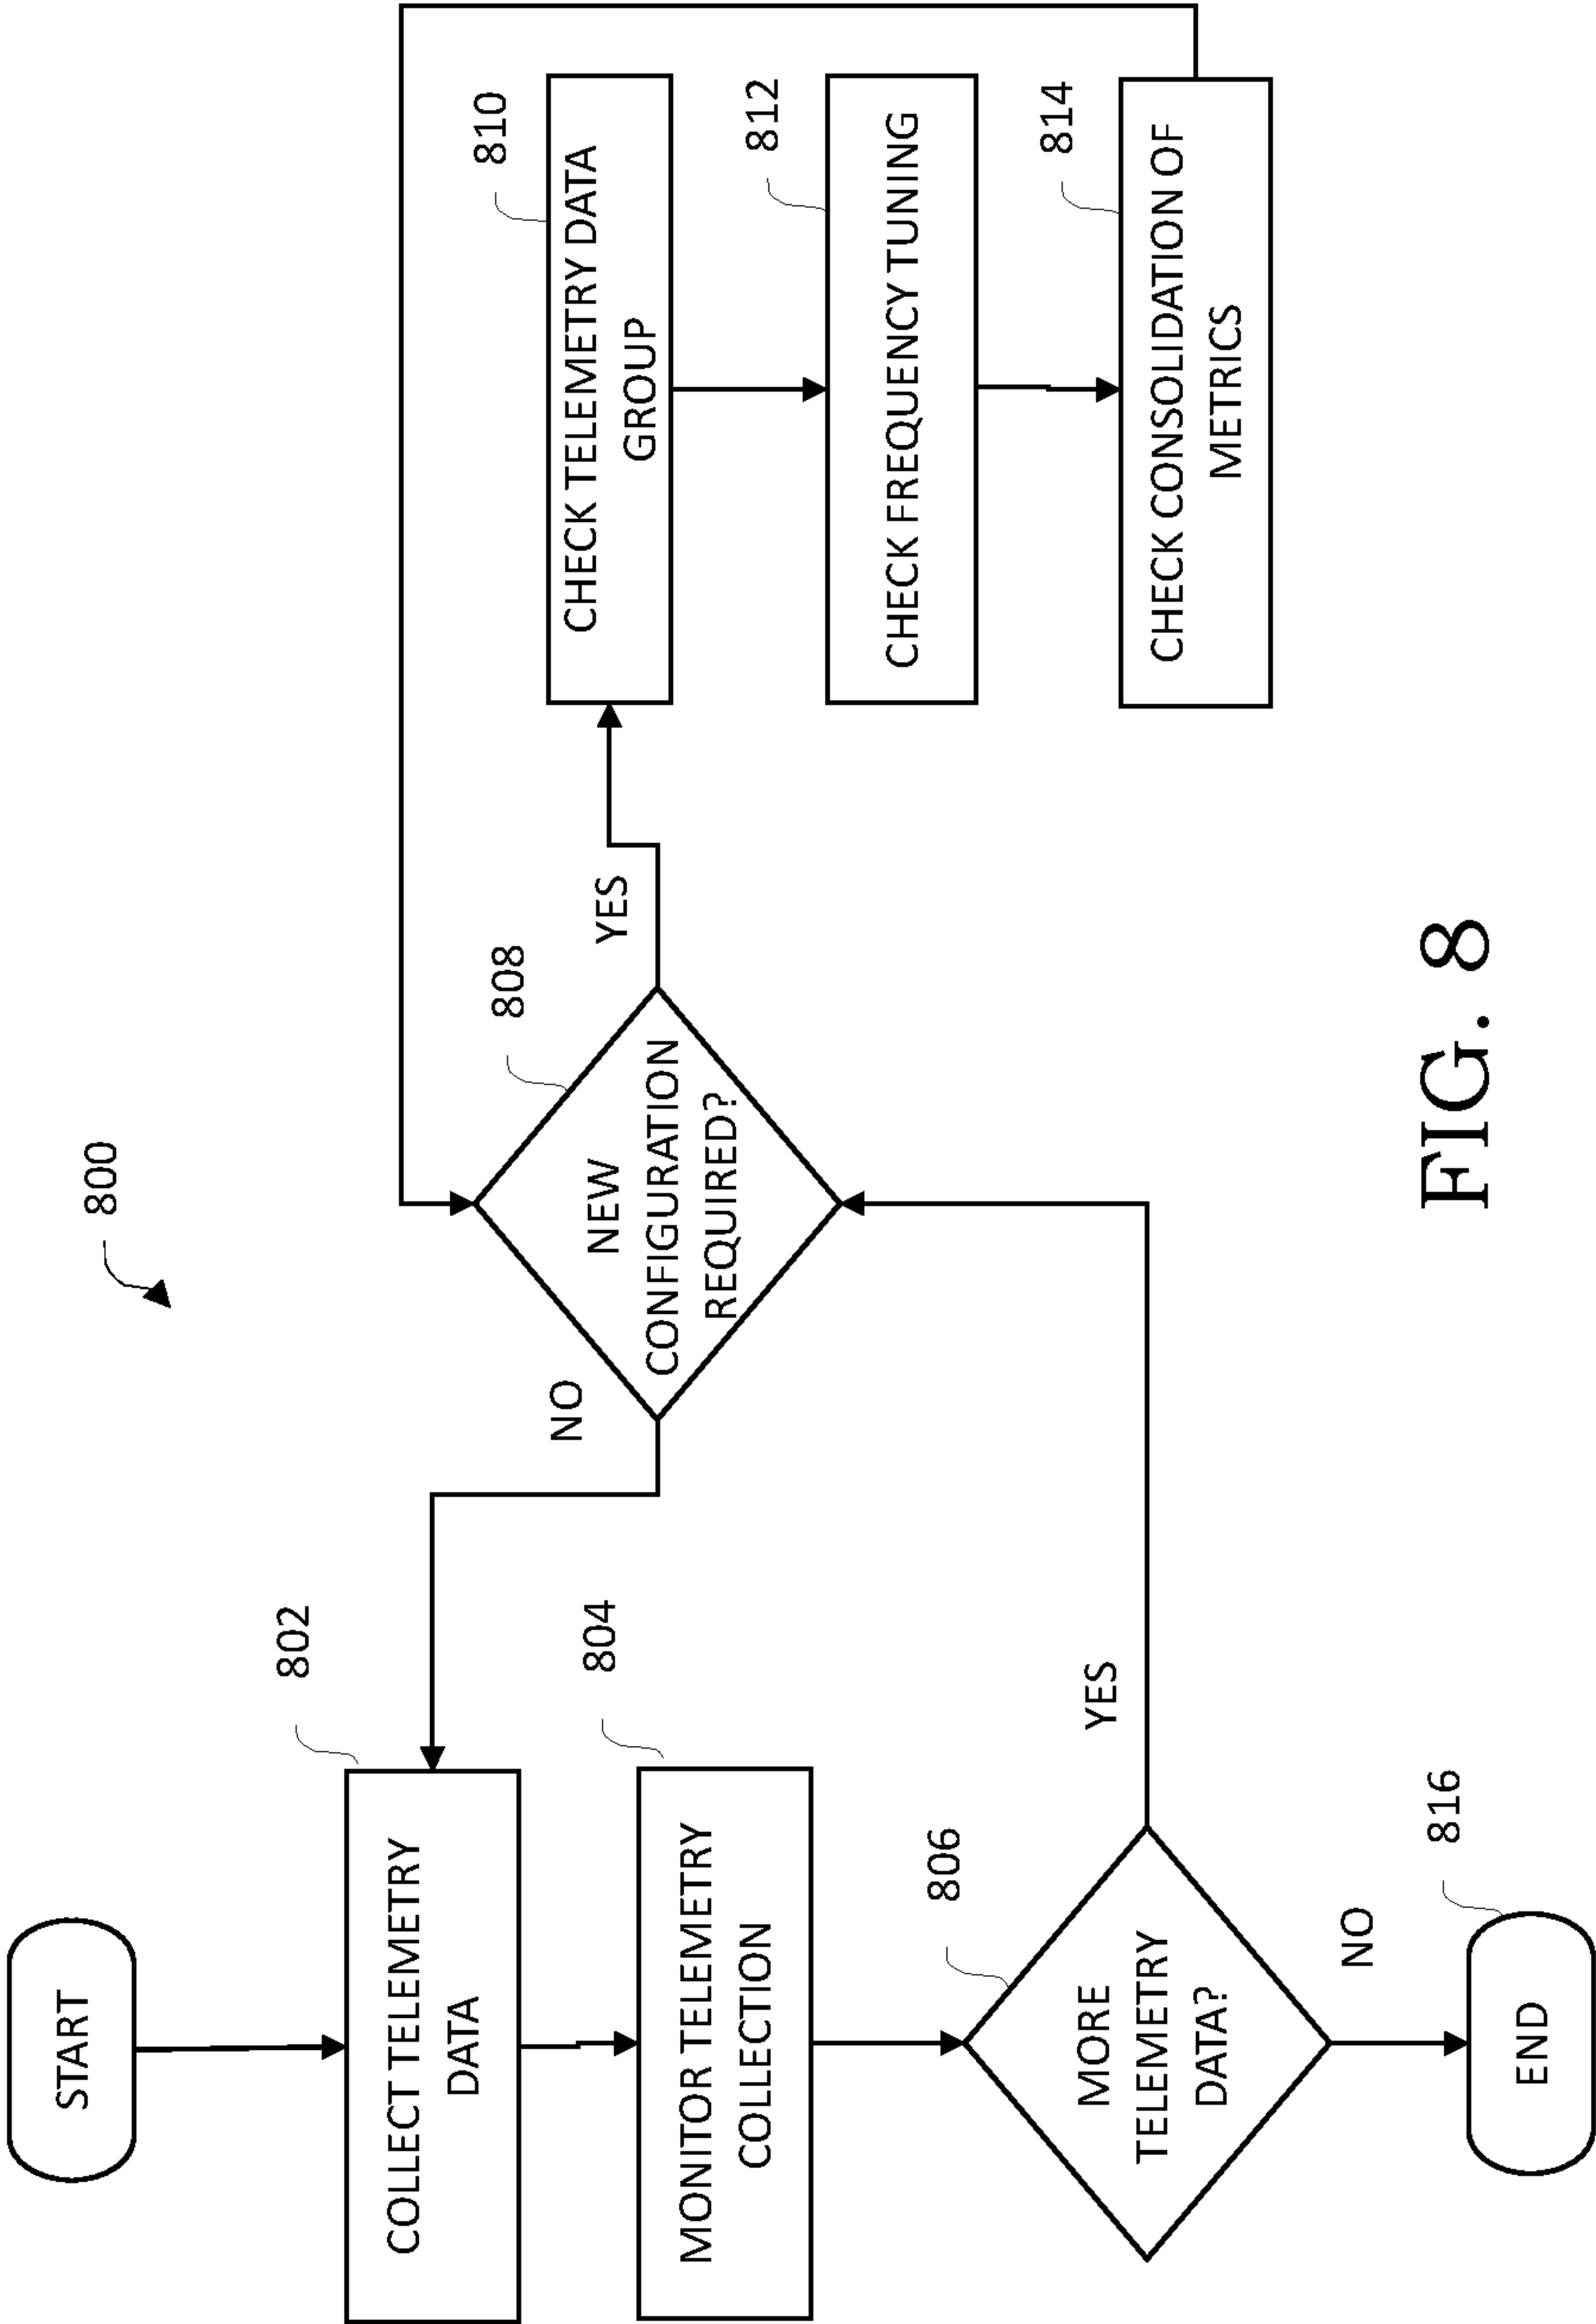


FIG. 8

## SYSTEM AND METHOD FOR TELEMETRY STREAMING AS A SERVICE

### BACKGROUND

#### Field of the Disclosure

**[0001]** This disclosure relates generally to information handling systems and more particularly to remote access controllers found in servers, and more particularly to systems and methods for providing remote access controller services.

#### Description of the Related Art

**[0002]** As the value and use of information continues to increase, individuals and businesses seek additional ways to process and store information. One option available to users is information handling systems. An information handling system generally processes, compiles, stores, and/or communicates information or data for business, personal, or other purposes thereby allowing users to take advantage of the value of the information. Because technology and information handling needs and requirements vary between different users or applications, information handling systems may also vary regarding what information is handled, how the information is handled, how much information is processed, stored, or communicated, and how quickly and efficiently the information may be processed, stored, or communicated. The variations in information handling systems allow for information handling systems to be general or configured for a specific user or specific use such as financial transaction processing, airline reservations, enterprise data storage, or global communications. In addition, information handling systems may include a variety of hardware and software components that may be configured to process, store, and communicate information and may include one or more computer systems, data storage systems, and networking systems.

### SUMMARY

**[0003]** Embodiments may be directed to a system for telemetry streaming, the system comprising a remote access controller (RAC) on a server and a private cloud server (PCS). The RAC may comprise a RAC processor and a RAC memory storing a set of instructions executable by the RAC processor to collect a first set of telemetry data associated with a plurality of devices on the server based on one or more telemetry collection requirements. The PCS may comprise a PCS processor and a PCS memory storing a set of available services and a set of instructions executable by the PCS processor to: communicate with the RAC to receive the first set of telemetry data from the RAC; communicate with the RAC to monitor outgoing telemetry traffic associated with the first set of telemetry data; store the first set of telemetry data in a telemetry database; determine the one or more telemetry collection requirements; and communicate a second set of telemetry data to an external cloud server, wherein the second set of telemetry data communicated to the external cloud server contains less data than the first set of telemetry data and meets the one or more telemetry collection requirements.

**[0004]** In some embodiments, the PCS memory stores a telemetry filter service configured to filter the first set of telemetry data. In some embodiments, the PCS memory

stores a telemetry aggregator service configured to aggregate the first set of telemetry data into the second set of telemetry data. In some embodiments, the PCS memory stores a telemetry streaming service configured to determine a frequency of communicating the second set of telemetry data to the external cloud server. In some embodiments, the set of PCS instructions are executable by the PCS processor for communicating the second set of telemetry data to the external cloud server based on the second set of telemetry data corresponding to a mandatory list of telemetry data. In some embodiments, the set of PCS instructions are executable by the PCS processor for communicating the second set of telemetry data to the external cloud server based on the second set of telemetry data corresponding to an error list of telemetry data. In some embodiments, the set of PCS instructions are executable by the PCS processor for communicating the second set of telemetry data to the external cloud server based on the second set of telemetry data corresponding to an optional list of telemetry data.

**[0005]** Embodiments may be directed to a data center comprising a plurality of servers, wherein each server comprises a remote access controller (RAC) comprising a RAC processor and a RAC memory storing a set of instructions executable by the RAC processor to collect a first set of telemetry data associated with the plurality of devices on the server based on one or more telemetry collection requirements. At least one server of the plurality of servers comprises a private cloud server (PCS) communicatively coupled to the RAC in each server of the plurality of servers and an external cloud server, the private cloud server comprising: a PCS processor; and a PCS memory storing a set of available services and a set of instructions executable by the PCS processor to: communicate with the RAC to receive the first set of telemetry data from the RAC in each server of the plurality of servers; communicate with the RAC in each server of the plurality of servers to monitor outgoing telemetry traffic associated with the first set of telemetry data; store the first set of telemetry data in a telemetry database; determine the one or more telemetry collection requirements; and communicate a second set of telemetry data to the external cloud server, wherein the second set of telemetry data communicated to the external cloud server contains less data than the first set of telemetry data and meets the one or more telemetry collection requirements.

**[0006]** In some embodiments, the PCS memory stores a telemetry filter service configured to filter the first set of telemetry data. In some embodiments, the PCS memory stores a telemetry aggregator service for aggregating the first set of telemetry data received from the plurality of servers into the second set of telemetry data. In some embodiments, the PCS memory stores a telemetry streaming service for determining a frequency of communicating the second set of telemetry data to the external cloud server. In some embodiments, the set of PCS instructions are executable by the PCS processor for communicating the second set of telemetry data to the external cloud server based on the second set of telemetry data corresponding to a mandatory list of telemetry data. In some embodiments, the set of PCS instructions are executable by the PCS processor for communicating the second set of telemetry data to the external cloud server based on the second set of telemetry data corresponding to an error list of telemetry data. In some embodiments, the set of PCS instructions are executable by the PCS processor for communicating the second set of telemetry data to the



external cloud server based on the second set of telemetry data corresponding to an optional list of telemetry data.

**[0007]** Embodiments may be directed to a method of operating a data center comprising a plurality of servers. The method comprises: storing, on each server of the plurality of servers, a set of instructions executable by the RAC processor to collect a first set of telemetry data associated with a plurality of devices on the server based on one or more telemetry collection requirements; executing, by a private cloud server communicatively coupled to the RAC in each server of the plurality of servers, a set of instructions to perform: communicating with the RAC in each server of the plurality of servers to receive the first set of telemetry data from the RAC in each server of the plurality of servers; communicating with the RAC in each server of the plurality of servers to monitor outgoing telemetry traffic associated with the first set of telemetry data; storing the first set of telemetry data received from the plurality of servers in a telemetry database; determining one or more telemetry collection requirements; and communicating a second set of telemetry data to an external cloud server, wherein the second set of telemetry data communicated to the external cloud server contains less data than the first set of telemetry data and meets the one or more telemetry collection requirements.

**[0008]** In some embodiments, the method comprises filtering the first set of telemetry data received from the plurality of servers into the second set of telemetry data. In some embodiments, the method comprises aggregating the first set of telemetry data received from the plurality of servers into the second set of telemetry data. In some embodiments, the method comprises a frequency of communicating the second set of telemetry data to the external cloud server. In some embodiments, the method comprises communicating the second set of telemetry data to the external cloud server based on the second set of telemetry data corresponding to one of a mandatory list of telemetry data, an error list of telemetry data or an optional list of telemetry data.

#### BRIEF DESCRIPTION OF THE DRAWINGS

**[0009]** For a more complete understanding of the present disclosure and its features/advantages, reference is now made to the following description, taken in conjunction with the accompanying drawings, which are not drawn to scale, and in which:

**[0010]** FIG. 1 depicts an example architecture of a plurality of data centers, with each data center containing a plurality of servers, with each server communicatively coupled to an external cloud server for managing the server;

**[0011]** FIG. 2 depicts an example of a server with a Remote Access Controller (RAC), illustrating a set of services commonly installed in firmware on a RAC for remotely monitoring and managing the server;

**[0012]** FIG. 3 depicts a system architecture of a plurality of data centers communicatively coupled to an external cloud server, with each data center containing a plurality of servers and a private cloud server, with each server being communicatively coupled to the private cloud server for managing the servers, in accordance with some embodiments;

**[0013]** FIG. 4 depicts a portion of the system architecture of FIG. 3, illustrating an embodiment of a RAC that may be

installed in a server and communicatively coupled to a private cloud server and one or more external cloud servers for supporting the server;

**[0014]** FIG. 5 depicts a portion of the system architecture of FIG. 3, illustrating one embodiment of a private cloud server with firmware installed and communicatively coupled to a RAC in a server and one or more external cloud servers;

**[0015]** FIG. 6 depicts one embodiment of a deployment architecture for providing available services to a RAC;

**[0016]** FIG. 7 depicts a portion of the system architecture of FIG. 3, illustrating one embodiment of a private cloud server with firmware installed and communicatively coupled to a RAC in a server and one or more external cloud servers; and

**[0017]** FIG. 8 depicts a flow diagram illustrating a method for configuring a private cloud server based on meeting telemetry data collection requirements for one or more client information handling systems.

#### DETAILED DESCRIPTION

**[0018]** In the following description, details are set forth by way of example to facilitate discussion of the disclosed subject matter. It should be apparent to a person of ordinary skill in the field, however, that the disclosed embodiments are examples and not exhaustive of all possible embodiments.

**[0019]** As used herein, a reference numeral refers to a class or type of entity, and any letter following such reference numeral refers to a specific instance of a particular entity of that class or type. Thus, for example, a hypothetical entity referenced by '12A' may refer to a particular instance of a particular class/type, and the reference '12' may refer to a collection of instances belonging to that particular class/type or any one instance of that class/type in general.

**[0020]** An information handling system (IHS) may include a hardware resource or an aggregate of hardware resources operable to compute, classify, process, transmit, receive, retrieve, originate, switch, store, display, manifest, detect, record, reproduce, handle, and/or utilize various forms of information, intelligence, or data for business, scientific, control, entertainment, or other purposes, according to one or more embodiments. For example, an IHS may be a personal computer, a desktop computer system, a laptop computer system, a server computer system, a mobile device, a tablet computing device, a personal digital assistant (PDA), a consumer electronic device, an electronic music player, an electronic camera, an electronic video player, a wireless access point, a network storage device, or another suitable device and may vary in size, shape, performance, functionality, and price. In one or more embodiments, a portable IHS may include or have a form factor of that of or similar to one or more of a laptop, a notebook, a telephone, a tablet, and a PDA, among others. For example, a portable IHS may be readily carried and/or transported by a user (e.g., a person). In one or more embodiments, components of an IHS may include one or more storage devices, one or more communications ports for communicating with external devices as well as various input and output (I/O) devices, such as a keyboard, a mouse, and a video display, among others. In one or more embodiments, IHS may include one or more buses operable to transmit communication between or among two or more hardware components. In one example, a bus of an IHS may include one or more of a memory bus, a peripheral bus, and a local



bus, among others. In another example, a bus of an IHS may include one or more of a Micro Channel Architecture (MCA) bus, an Industry Standard Architecture (ISA) bus, an Enhanced ISA (EISA) bus, a Peripheral Component Interconnect (PCI) bus, HyperTransport (HT) bus, an inter-integrated circuit (I<sup>2</sup>C) bus, a serial peripheral interface (SPI) bus, a low pin count (LPC) bus, an enhanced serial peripheral interface (eSPI) bus, a universal serial bus (USB), a system management bus (SMBus), and a Video Electronics Standards Association (VESA) local bus, among others.

**[0021]** In one or more embodiments, an IHS may include firmware that controls and/or communicates with one or more hard drives, network circuitry, one or more memory devices, one or more I/O devices, and/or one or more other peripheral devices. For example, firmware may include software embedded in an IHS component utilized to perform tasks. In one or more embodiments, firmware may be stored in non-volatile memory, such as storage that does not lose stored data upon loss of power. In one example, firmware associated with an IHS component may be stored in non-volatile memory that is accessible to one or more IHS components. In another example, firmware associated with an IHS component may be stored in non-volatile memory that may be dedicated to and includes part of that component. For instance, an embedded controller may include firmware that may be stored via non-volatile memory that may be dedicated to and includes part of the embedded controller.

**[0022]** An IHS may include a processor, a volatile memory medium, non-volatile memory media, an I/O subsystem, and a network interface. Volatile memory medium, non-volatile memory media, I/O subsystem, and network interface may be communicatively coupled to processor. In one or more embodiments, one or more of volatile memory medium, non-volatile memory media, I/O subsystem, and network interface may be communicatively coupled to processor via one or more buses, one or more switches, and/or one or more root complexes, among others. In one example, one or more of a volatile memory medium, non-volatile memory media, an I/O subsystem and network interface may be communicatively coupled to the processor via one or more PCI-Express (PCIe) root complexes. In another example, one or more of an I/O subsystem and a network interface may be communicatively coupled to processor via one or more PCIe switches.

**[0023]** In one or more embodiments, the term “memory medium” may mean a “storage device”, a “memory”, a “memory device”, a “tangible computer readable storage medium”, and/or a “computer-readable medium”. For example, computer-readable media may include, without limitation, storage media such as a direct access storage device (e.g., a hard disk drive, a floppy disk, etc.), a sequential access storage device (e.g., a tape disk drive), a compact disk (CD), a CD-ROM, a digital versatile disc (DVD), a random access memory (RAM), a read-only memory (ROM), a one-time programmable (OTP) memory, an electrically erasable programmable read-only memory (EEPROM), and/or a flash memory, a solid state drive (SSD), or any combination of the foregoing, among others.

**[0024]** In one or more embodiments, one or more protocols may be utilized in transferring data to and/or from a memory medium. For example, the one or more protocols may include one or more of small computer system interface (SCSI), Serial Attached SCSI (SAS) or another transport

that operates with the SCSI protocol, advanced technology attachment (ATA), serial ATA (SATA), a USB interface, an Institute of Electrical and Electronics Engineers (IEEE) 1394 interface, a Thunderbolt interface, an advanced technology attachment packet interface (ATAPI), serial storage architecture (SSA), integrated drive electronics (IDE), or any combination thereof, among others.

**[0025]** A volatile memory medium may include volatile storage such as, for example, RAM, DRAM (dynamic RAM), EDO RAM (extended data out RAM), SRAM (static RAM), etc. One or more of non-volatile memory media may include nonvolatile storage such as, for example, a read only memory (ROM), a programmable ROM (PROM), an erasable PROM (EPROM), an electrically erasable PROM, NVRAM (non-volatile RAM), ferroelectric RAM (FRAM), a magnetic medium (e.g., a hard drive, a floppy disk, a magnetic tape, etc.), optical storage (e.g., a CD, a DVD, a BLU-RAY disc, etc.), flash memory, a SSD, etc. In one or more embodiments, a memory medium can include one or more volatile storages and/or one or more nonvolatile storages.

**[0026]** In one or more embodiments, a network interface may be utilized in communicating with one or more networks and/or one or more other information handling systems. In one example, network interface may enable an IHS to communicate via a network utilizing a suitable transmission protocol and/or standard. In a second example, a network interface may be coupled to a wired network. In a third example, a network interface may be coupled to an optical network. In another example, a network interface may be coupled to a wireless network. In one instance, the wireless network may include a cellular telephone network. In a second instance, the wireless network may include a satellite telephone network. In another instance, the wireless network may include a wireless Ethernet network (e.g., a Wi-Fi network, an IEEE 802.11 network, etc.).

**[0027]** In one or more embodiments, a network interface may be communicatively coupled via a network to a network storage resource. For example, the network may be implemented as, or may be a part of, a storage area network (SAN), personal area network (PAN), local area network (LAN), a metropolitan area network (MAN), a wide area network (WAN), a wireless local area network (WLAN), a virtual private network (VPN), an intranet, an Internet or another appropriate architecture or system that facilitates the communication of signals, data and/or messages (generally referred to as data). For instance, the network may transmit data utilizing a desired storage and/or communication protocol, including one or more of Fibre Channel, Frame Relay, Asynchronous Transfer Mode (ATM), Internet protocol (IP), other packet-based protocol, Internet SCSI (iSCSI), or any combination thereof, among others.

**[0028]** In one or more embodiments, a processor may execute processor instructions in implementing at least a portion of one or more systems, at least a portion of one or more flowcharts, at least a portion of one or more methods, and/or at least a portion of one or more processes. In one example, a processor may execute processor instructions from one or more memory media in implementing at least a portion of one or more systems, at least a portion of one or more flowcharts, at least a portion of one or more methods, and/or at least a portion of one or more processes. In another example, a processor may execute processor instructions via a network interface in implementing at least a portion of one



or more systems, at least a portion of one or more flowcharts, at least a portion of one or more methods, and/or at least a portion of one or more processes.

[0029] In one or more embodiments, a processor may include one or more of a system, a device, and an apparatus operable to interpret and/or execute program instructions and/or process data, among others, and may include one or more of a microprocessor, a microcontroller, a digital signal processor (DSP), an application specific integrated circuit (ASIC), and another digital or analog circuitry configured to interpret and/or execute program instructions and/or process data, among others. In one example, a processor may interpret and/or execute program instructions and/or process data stored locally (e.g., via memory media and/or another component of an IHS). In another example, a processor may interpret and/or execute program instructions and/or process data stored remotely (e.g., via a network storage resource).

[0030] In one or more embodiments, an I/O subsystem may represent a variety of communication interfaces, graphics interfaces, video interfaces, user input interfaces, and/or peripheral interfaces, among others. For example, an I/O subsystem may include one or more of a touch panel and a display adapter, among others. For instance, a touch panel may include circuitry that enables touch functionality in conjunction with a display that is driven by a display adapter.

[0031] A non-volatile memory medium may include an operating system (OS) and applications (APPs). In one or more embodiments, one or more of an OS and APPs may include processor instructions executable by a processor. In one example, a processor may execute processor instructions of one or more of OS and APPs via a non-volatile memory medium. In another example, one or more portions of the processor instructions of one or more of an OS and APPs may be transferred to a volatile memory medium and a processor may execute the one or more portions of the processor instructions.

[0032] Non-volatile memory medium may include information handling system firmware (IHSFW). In one or more embodiments, IHSFW may include processor instructions executable by a processor. For example, IHSFW may include one or more structures and/or one or more functionalities of and/or compliant with one or more of a basic input/output system (BIOS), an Extensible Firmware Interface (EFI), a Unified Extensible Firmware Interface (UEFI), and an Advanced Configuration and Power Interface (ACPI), among others. In one instance, a processor may execute processor instructions of IHSFW via non-volatile memory medium. In another instance, one or more portions of the processor instructions of IHSFW may be transferred to volatile memory medium, and processor may execute the one or more portions of the processor instructions of IHSFW via volatile memory medium.

[0033] Data centers may have large numbers of information handling systems such as servers for processing information. A data center facility may have one or more floors with each floor having racks of servers. A server may be processing a set of information independently or a group of servers may be working on the same set of information.

[0034] Turning now to FIG. 1, an architectural diagram of a plurality of data centers 10 illustrates that each data center 10 may contain a plurality of information handling systems 12 (also referred to as servers 12). Data center 10-1 may contain M servers 12 and data center 10-2 may contain N

servers, where each of M and N may be any number between two to over a thousand. Each server 12 may comprise Remote Access Controller (RAC) 14 to allow remote monitoring and management of any server 12 in data center 10. RAC 14 may comprise a Dell Remote Access Controller (DRAC) or an integrated Dell Remote Access Controller (iDRAC) for remote monitoring and management.

[0035] Turning to FIG. 2, each server 12 has remote access controller (RAC) 14 with firmware to allow for remote monitoring and managing of server 12, and server 12 may need to communicate with one or more other servers 12 in data center 10. Firmware within RAC 14 may include a Representational State Transfer Application Programming Interface (REST API) 16 that allows servers 12 to communicate with each other. Remote monitoring may require communications with processes running in RAC 14. A D-Bus interface 18 or other inter-process communications mechanism may enable local communication between processes stored in firmware on RAC 14.

[0036] Firmware in RAC 14 may include a set of built in controls 20 and features 30 executable by a processor in RAC 14. For example, each server 12 has a power supply unit (PSU) for receiving power from a power source and transforming the power into usable power by the server 12. Firmware in RAC 14 includes power control 20-1 for monitoring of power received by the PSU to determine when server 12 has switched power modes and allow for remote switching of server 12 between power modes. Similarly, server 12 has a fan for cooling server 12 and firmware in RAC 14 includes thermal control 20-2 for monitoring of temperatures in server 12 to determine if the fan is operating, determining when to operate the fan and enabling remote operation of the fan.

[0037] Firmware in RAC 14 may also store hardware abstraction layer (HAL) 22 for communication with and monitoring of peripheral devices. Within HAL 22, features such as Inter-Integrated Circuit (I2C) protocol 22-1, Improved Inter-Integrated Circuit (I3C) protocol 22-2, Serial Peripheral Interface (SPI) 22-3, Reduced Media-Independent Interface (RMII) 22-4, Peripheral Connect Interface Express Vendor Defined Message (PCIe VDM) 22-5 and management component transport protocol (MCTP) 22-6 allow communications with chips, processors and microcontrollers in server 12.

[0038] Server 12 includes hardware 24 such as processors for processing information, memory for storing information, a fan for cooling server 12 including devices and components in server 12, a network integrated circuit (NIC) for communications, other controllers such as RAID controllers, and Complex Logic Programmable Devices (CPLDs). Accordingly, Firmware in RAC 14 may include other features such as Basic Input Output Service (BIOS) 26-1, NIC 26-2, fan control 26-3, PSU control 26-4 for operating a PSU, RAID feature 26-5 for managing a RAID and CPLD feature 26-6 for monitoring and managing a CPLD.

[0039] Firmware in RAC 14 may typically include other features 30 for monitoring and managing servers 12 in data centers 10. The examples depicted in FIG. 3 and provided below represent more common features found in RAC 14, and RAC 14 may store additional features 30 not listed.

[0040] Web Graphical User Interface (GUI) 30-1 is a web-based application that processes network events for displaying the data in a graphical format for users to view on an attached display. Redfish® 30-2 is an application pro-



gramming interface (API) that uses RESTful semantics to access data defined in model format for systems management. Remote Access Controller admin (RACADM) feature **30-3** is a scriptable interface to allow remote configuration of RAC **14**. Simple Network Management Protocol (SNMP) feature **30-4** may be used to collect data related to network changes and determine the status of network-connected devices. Hardware inventory feature **30-5** may maintain an inventory and properties of all hardware installed on server **12**. Software inventory feature **30-6** may maintain an inventory and versions of all software running on server **12**. Firmware update feature **30-7** may maintain a list of firmware including versions and facilitate updates of any firmware on server **12**. Monitoring feature **30-8** may monitor operation of components or devices in server **12** and record values or may define what operations are to be monitored and how the monitoring should occur. RAC **14** may include database **32** for storing information about components or devices in server **12**. Network Controller-Sideband Interface (e.g., NC-SI 1.2) feature **30-10** defines a control communication protocol between a baseboard management controller (BMC) and one or more Network Interface Controllers (NICs). Platform Level Data Model (PLDM) feature **30-11** defines the contents of a firmware update package. Security Protocol and Data Models (SPDM) feature **30-12** enables authentication, attestation and key exchange for enabling security of server **12**. Management Control Transport Protocol (MCTP) feature **30-13** stores message formats, transport descriptions, message exchange patterns and endpoint characteristics related to communications between components.

[0041] Servers **12** described above have several shortcomings. All features **30** are stored in memory in each RAC **14** of each server **12**, regardless of whether a feature **30** will be used for a particular server **12**. Furthermore, features **30** are getting more robust and require more memory to manage devices in servers **12**. For example, RAC **14** in some generations of servers **12** may have 512 Megabytes (MB) of memory to store all features **30**, RAC **14** in later generations of servers **12** may have 1 Gigabyte (GB) of memory to store all features **30**, and RAC **14** in later generations of servers **12** may have 2 GB of memory to store all features **30**. Although each server **12** is provided with the full suite of features **30**, many customers or servers **12** require a limited set of features **30**. For security reasons, some users prefer to have only firmware code which they know they will use. Subscription-based licensing, in which a user pays for only the features needed, may be preferred. However, only about 35% of datacenter servers **12** contain the latest firmware for RAC **14**, and features **30** are often unnecessarily tied to specific hardware capabilities. Customization of features **30** via Software Developer Kits (SDK) may be limited or unavailable due to hardware capabilities or software requirements. Some devices (e.g., Baseboard Management Controllers (BMCs)) are expected to be commoditized, resulting in greater complexity and/or more features **30** that would need to be installed on each RAC **14**.

[0042] Turning to FIGS. 3, 4 and 5, embodiments disclosed herein allow a remote access controller (RAC) **314** in server **312** to have firmware that has a minimum number of features **30** that can be augmented as needed with any number of available services **350** executing on a private cloud server **360**. In some embodiments, RAC **314** may be manufactured with firmware including a minimum Hard-

ware Abstraction Layer (HAL). Some controls **20** and features **30** may remain as built-in controls **20** or features **30** in firmware in RAC **314**. For example, any time server **12** is operating, PSU control **20-1** monitors power supplied to the PSU and is therefore maintained in firmware in RAC **314**. Embodiments may augment built-in features **30** with available services **350** that can be retrieved from external cloud server **210** in cloud **4** and installed and executed on private cloud server **360** or possibly installed on RAC **314**, discussed in greater detail below. If needed by RAC **314**, available services **350** may be delivered/deployed by leveraging and extending on top of existing firmware in RAC **314**.

[0043] Advantageously, embodiments may provide vendor-agnostic, unified and seamless system management across generations of servers **312**. Embodiments may further provide accelerated delivery of new available services **350** with subscription-based granular licensing.

[0044] Referring to FIG. 3, data centers **310** may each have multiple servers **312**. Data center **310-1** may have a first plurality of servers **312** (e.g., 2 to M) divided into P groups **340**, wherein group **340-1** may have a first set of servers **312** (e.g., 1 to A), group **340-2** may have a second set of servers **312** (e.g., 1 to B) and group **340-P** may have a third set of servers **312** (e.g., 1 to C). Data center **310-1** may further comprise private cloud server **360-1** communicatively coupled to all servers **312-1** to **312-M** in data center **310-1**. Private cloud server **360-1** may be communicatively coupled to external cloud server **210** in cloud **4**.

[0045] Data center **310-2** may have a second plurality of servers **312** (e.g., 2 to N) divided into Q groups **340**, wherein group **340-4** may have a first set of servers **312** (e.g., 1 to D), group **340-2** may have a second set of servers **312** (e.g., 1 to E) and group **340-Q** may have a third set of servers **312** (e.g., 1 to F). Data center **310-2** may further comprise private cloud server **360-2** communicatively coupled to all servers **312-1** to **312-N** in data center **310-2**. Private cloud server **360-2** may be communicatively coupled to external cloud servers **210** in cloud **4**.

[0046] Groups **340** may be logical groups, discussed in greater detail with respect to FIG. 6.

[0047] Referring to FIG. 4, external cloud server **210** may communicate with private cloud servers **360** over a network to provide available services **350**.

[0048] Cloud integration service **210-1** may include data monitoring **214** and cloud services **216**. Dell CloudIQ is an example of a data monitoring service **214** that may monitor operation of servers **312** in data centers **310** for data protection. Dell APEX is an example catalog of cloud services **216** that may integrate infrastructure and software for cloud operations.

[0049] Infrastructure Solutions Group (ISG) server **210-2** may provide storage and networking services and may include third-party services **350-18**. In some embodiments, ISG server **210-2** may provide RAC proxy on Data Processing Unit (DPU) service **220** to private cloud server **360** to offload processing.

[0050] Management Service Catalog (MSC) server **210-3** may store a database of available services **350** that can be provided to private cloud servers **360**. In some embodiments, private cloud server **360** communicates with MSC server **210-3** to retrieve available services **350**. In some embodiments, MSC server **210-3** communicates with private cloud server **360** to install available services **350**.



[0051] IoT integration service **210-4** may enable devices to communicate with other devices and enable external cloud servers **210** to communicate with or monitor devices in servers **312**.

[0052] Autonomous server manager **210-5** may provide server management without a user.

[0053] Turning to one or more of FIGS. **4** and **5**, servers **312** may be similar to servers **12** depicted in FIG. **2** in that they comprise hardware such as a PSU for receiving power from a source and transforming the power into usable power by the server, a fan, network IC (NIC) for communications, other controllers such as RAID controllers, and Complex Logic Programmable Devices (CPLDs). Servers **312** also include BIOS **26-1**, NIC **26-2**, fan control **26-3**, PSU service **26-4** for operating a PSU, remote controller service **26-5** for managing a RAID and CPLD service **26-6** for monitoring and managing CPLDs.

[0054] Firmware within RAC **314** on servers **312** may include a Representational State Transfer Application Programming Interface (REST API) **16** that allows servers **312** to communicate with each other. Remote monitoring may require communications with processes running in RAC **314**. A D-Bus **18** or other inter-process communications mechanism enables local communication between processes on RAC **314**. RAC **314** may further comprise Remote Procedure Call (RPC) queue **302**, D-Bus to RPC mapper **304** and D-Bus to REST mapper **306** for available services **350** on private cloud server **360** to communicate with HAL **22**. In some embodiments, RPC mapper **304** may be a Google RPC (gRPC) mapper **304**.

[0055] RAC **314** may also store hardware abstraction layer (HAL) **22** for communication and monitoring of peripheral devices. Within HAL **22**, **12C 22-1**, **13C 22-2**, **SPI 22-3**, **RMII 22-4**, **PCIe VDM 22-5** and **MCTP 22-6** allow communications with chips, processors and microcontrollers in server **312**.

#### Available Services May be Stored and Executed on a Private Cloud Server

[0056] Embodiments of private cloud server **360** may store a set of available services **350** retrieved from MSC server **210-3**. In some embodiments, private cloud server **360** may determine, based on one or more of an application to be executed on a server **312** or a data set to be processed by a server **312**, that a particular available service is needed and communicate with MSC server **210-3** to retrieve the available service **350**. In some embodiments, private cloud server **360** may communicate with MSC server **210-3** to retrieve an available service **350** based on a subscription. Available services **350** may be executed by private cloud server **360** unless there is a need to have RAC **314** execute the available service **350**.

#### Available Services May be Installed in a Rac if Needed

[0057] Embodiments may install available services **350** in RAC **314** from private cloud server **360**, wherein the set of available services **350** stored in RAC **314** may be less than the total number of available services **350** stored in private cloud server **360** and less than the plurality of available services **350** stored in MSC server **210-3**. Available services **350** may be installed in RAC **314** based on one or more of performance or capability. Once the process is complete or

the available service **350** is no longer time-sensitive, embodiments may uninstall the available service **350** from RAC **314**.

[0058] As a device-based example, referring to available service **350-13** (i.e., MCTP **350-13**) and available service **350-10** (i.e., NC-SI 1.2 protocols **350-10**), if a process sends a request to CPLD **26-6**, a response may be expected in a minimum response time of about 50 milliseconds (ms) and a maximum response time of about 4 seconds. If the MCTP **350-13** and NC-SI 1.2 protocols **350-10** are stored in private cloud server **360**, a response may take 5 seconds. Embodiments may install NC-SI 1.2 protocols **350-10** as a time sensitive plug-in **340** to ensure a response is received in less than 4 seconds. Once NC-SI 1.2 protocols **350-10** is not needed, NC-SI 1.2 protocols **350-10** may be uninstalled from RAC **314**.

[0059] As another example, referring to available service **350-11** (i.e., PLDM **350-11**) and available service **350-7** (i.e., Firmware update service **350-7**), if there is a firmware update and PLDM **350-11** executing on private cloud server **360** cannot deliver the update payload within a maximum time (e.g., 4 seconds), the firmware update may time out. In this case, PLDM **350-11** and Firmware update service **350-7** may be installed as time sensitive plug-ins **340** to ensure the firmware update payload can be delivered in time. Once the firmware update payload is delivered, PLDM **350-11** and Firmware update service **350-7** may be uninstalled from RAC **314**.

#### Existing Available Services Stored on Private Cloud Server

[0060] Box **352** contains a non-limiting list of available services **350** that may be installed on private cloud server **360** to communicate with RACs **314** on any servers **312** in data centers **310**, wherein available services **350** in box **352** may function similar to features **30** described above with respect to FIG. **2** but execute on private cloud server **360**. For example, RAC **314** may be configured with a minimum number of features **30** and private cloud server **360** may store and execute Web GUI service **350-1**; Redfish® server management service **350-2**; RACADM command-line service **250-3**; SNMP service **350-4**; Hardware (HW) inventory service **350-5** for tracking information and properties of hardware installed on servers **312**; Firmware (FW) inventory service **350-6** for tracking information and versions of firmware installed on servers **312**; FW update service **350-7** for installing Firmware updates in servers **312**; Monitoring service **350-8** for coordinating monitoring of operation of servers **312**; Unified database **350-9** for storing information for sharing among processes and applications executing on servers **312**; NC-SI service **350-10**; PLDM service **350-11**; SPDM service **350-12** and MCTP service **350-13**. In some embodiments, available services **350** may be retrieved from Management Service Catalog server **210-3** through subscriptions, wherein any available service **350** may be retrieved as needed.

#### New Services and Services with Increased Scope

[0061] Embodiments may allow private cloud server **360** to provide additional available services **350** and available services **350** with increased scope that increase the capabilities of RAC **314**. Box **354** contains a non-limiting list of available services **350** that may be retrieved from MSC server **210-3** and installed on private cloud server **360**. Some available services **350** in box **354** may be executed by



private cloud server **360**. Some available services **350** in box **354** may be installed on RAC **314** on any server **312** in data center **310**.

#### Credentials Service

[0062] A user must have administrative credentials to run RACADM commands remotely. When a user wants to run RACADM commands, the user must first be validated by RAC **314**. Traditionally, RAC **14** corresponding to server **12** in FIG. **2** may store credentials for one to about sixteen users. Embodiments described herein allow private cloud server **360** to download and execute credentials service **350-15**, which may be configured to manage credentials for one user to more than sixteen users and map the credentials to RACs **314** as needed. In some embodiments, unified database **350-9** on private cloud server **360** may be configured to store credentials for a large number of users (e.g., over a thousand) to increase the scope of credentials service **350-9** over credentialing features **30** commonly found in RACs **14**. In some embodiments, private cloud server **360** may be configured to install credential service **350-15** and store a database in unified database **350-9** with the credentials for any number of users between one to over a thousand. As the number of users changes, private cloud server **360** may update the database. In some embodiments, credentials service **350-15** may validate users. In some embodiments, credentials service **350-15** may register itself as a credentials manager, receive requests, and validate users such that RAC **314** does not validate but knows the user is validated.

#### Telemetry Service

[0063] Some Baseboard Management Controllers (BMCs) have no intelligence to adjust telemetry collection based on server events/errors or have different streaming rate for each metric. Traditionally, RAC **14** in server **12** may send the same information multiple times or at different rates, tying up networking resources and memory. Telemetry service **350-16** may be retrieved from MSC server **210-3** and installed on private cloud server **360**. In some embodiments, an Open Telemetry (OTEL) service **350-16** may be stored in private cloud server **360** as an available service **350**. When servers **312** communicate data for telemetry purposes, telemetry service **350-16** executing on private cloud server **360** may aggregate data, remove redundant data, or otherwise coordinate the communication of data, resulting in reduced network congestion, described in greater detail below. In some embodiments, telemetry service **350-16** may be installed in RAC **314** to meet telemetry requirements and then uninstalled after server **312** does not need to meet any telemetry requirements. In some embodiments, telemetry service **350-16** may be installed in RAC **314** as a time-sensitive plug-in **340** to provide quicker responses to telemetry requirements and then uninstalled after server **312** does not need quick responses to meet telemetry requirements.

#### AI/ML Service

[0064] Artificial Intelligence (AI)/Machine learning (ML) service **350-17** may include services necessary for AI/ML. If a server **312** (or a set of servers **312**) is needed for AI/ML, AI/ML service **350-17** may be downloaded to private cloud server **360** for coordinating processing by servers **312** for AI/ML processing.

#### Third Party Services

[0065] Third party services **350-18** may include services needed for particular third-party applications. Advantageously, instead of all available services **350** being tied to particular hardware, embodiments may enable third-party services **350-18** to execute on private cloud server **360**, wherein other available services **350**, RPC queue **302**, D-Bus to RPC mapper **304** and D-Bus to REST mapper **306** enable third-party services **350-18** on private cloud server **360** to communicate with HAL **22** in one or more RACs **314**.

#### Advanced Power/Thermal Service

[0066] Advanced power/thermal service **350-19** may be stored on private cloud server **360** and may refer to an available service **350** that can be executed on private cloud server **360** to communicate with RACs **314** on multiple servers **312** to monitor or control power or thermal operations of one or more servers **312**. For example, data center **310** may have multiple floors with hundreds of racks of servers **312**. Each RAC **314** may communicate with sensors inside a server **312** for remotely and individually monitoring temperature of that server **312**. Advanced power/thermal service **350-19** may allow a user to remotely and collectively monitor and manage temperatures for multiple servers **312**, such as all servers **312** processing a set of information, all servers **312** in a rack or all servers **312** on a floor. In some embodiments, advanced power/thermal service **350-19** may be installed in RAC **314** (e.g., as a time-sensitive plug-in **340**) for quicker response to power/thermal requirements and then uninstalled after server **312** does not need to operate under advanced power/thermal requirements.

#### Secure Zero Touch Provisioning (sZTP) Service

[0067] Secure Zero-Touch Provisioning (sZTP) service **350-20** enables a server **312** to securely provision a network device when booting in a default state. If server **312** is expected to require communication over a network, sZTP service **350-20** executing on private cloud server **360** may ensure the network device is securely provisioned. In some embodiments, sZTP service **350-20** may be installed in RAC **314** for booting and uninstalled once server **312** has successfully booted. In some embodiments, sZTP service **350-20** may be installed on private cloud server **360** and installed on servers **312** (e.g., as a time-sensitive plug-in **340**) as needed to ensure network devices are securely provisioned.

#### ISM Services

[0068] Integrated Dell Remote Access Controller (iDRAC) Service Manager service **350-21** may refer to available services **350-21** that may be executed to monitor and manage operation of RAC **314**.

#### Fido Device Onboard (FDO) Service

[0069] Fast IDentity Online (FIDO) Device Onboard (FDO) service **350-22** allows onboarding and trusting of new devices (e.g., RAC **314**) within an environment. For example, RAC **314** may be Linux-based and one or more available services **350** may be Windows-based. FDO service **350-22** may enable servers **312** to execute Linux-based services **350** and Windows-based services **350**. Data centers **310** may have hundreds or thousands of servers **312**. Devices in servers **312** may be removed and exchanged for newer devices. Instead of a remote user configuring each



device and deploying required applications, FDO service **350-22** allows multiple devices in various servers **312** to be configured correctly. In some embodiments, FDO service **350-22** may be installed in RAC **314** for quicker configuring and deployment and uninstalled once server **312** has successfully configured and deployed new devices. In some embodiments, FDO service **350-22** may be installed on private cloud server **360** and installed on servers **312** (e.g., as a time-sensitive plug-in **340**) as needed to ensure new devices are quickly configured and deployed.

#### Internet of Things (IOT) Integration Service

[0070] IoT Integration service **350-23** may be installed on private cloud server **360** to facilitate integration between devices on servers **312**. Advantageously, instead of each server **312** in multiple data centers **310** communicating with IoT integration server **210-4**, private cloud server **360** may perform some of the integration, wherein other available services **350**, RPC queue **302**, D-Bus to RPC mapper **304** and D-Bus to REST mapper **306** enable private cloud server **360** to communicate with HAL **22** in one or more RACs **314** in servers **312** that may have different devices.

#### Feature Tour

[0071] Feature Tour service **350-24** may be installed to guide users through features **30** available to server **312**. Advantageously, instead of each RAC **314** in each server **312** in multiple data centers **310** storing all the information necessary to guide users through only features **30** installed on that server, private cloud server **360** may store all the information and provide information including features **30** and available services **350** available to a particular server **312**. An end user may not know whether a feature is installed on server **312** or is instead an available service **350** accessible by server **312**. In some embodiments, private cloud server **360** may install feature tour service **350-24** with information relevant to a particular server **312** on RAC **314** associated with the particular server **312**.

#### Infrastructure Programmer Developer Kit (IPDK) Service

[0072] IPDK service **350-25** is an open-source, vendor-agnostic framework of drivers and APIs. IPDK service **350-25** may be too large to install on each RAC **314** in each server **312**. IPDK service **350-25** may be installed on private cloud server **360** for infrastructure offload and management.

[0073] For each available service **350**, if the available service **350** is needed for a particular server **312**, a version of the available service **350** may be retrieved from external cloud server **210** and executed on private cloud server **360**, wherein RAC **314** contains a minimum number of built-in controls **20** and features **30** to communicate with private cloud server **360**, wherein private cloud server **360** comprises memory and processors for executing the available service **350**. If, at some later time, the available service **350** is not needed on a particular server **312**, the available service **350** may be uninstalled from RAC **314** but a version may still be stored in private cloud server **360** or the available service **350** may be uninstalled from private cloud server **360**.

#### MSO—Management Service Orchestrator

[0074] MSO **356** may coordinate between features and ensure any service **350** is compatible with other available

services **350** and hardware in server **312** and that no conflicts exist. For example, regarding telemetry, a manufacturer may have a default telemetry service **20** installed in firmware, but a customer may want to use Open Telemetry (OTel) service **350-16** or a third-party telemetry service **350-18**. If the customer requests another telemetry service, MSO **356** may determine whether the requested telemetry application will work if an existing telemetry feature needs to be disabled or an existing telemetry service **350-16** or **350-18** needs to be uninstalled.

#### Large Send Offload (LSO)

[0075] Large Send Offload (LSO) service **358** may increase the egress throughput of high-bandwidth network connections. LSO service **358** may be installed on PCS **312**, reducing workload performed by RAC **314**. In some embodiments, LSO **358** enables communication with RAC Proxy on DPU service **220** to offload large workloads.

#### Provisioning Available Services with Redundancy

[0076] Referring to FIG. 6, embodiments may provide available services **350** to multiple RACs **314**. Each server **312** may have multiple firmware **602** associated with a Baseboard Management Controller (BMC) or a Remote Access Controller (RAC) **314**. For example, firmware **602-1** and **602-2** may correspond to firmware X and firmware Y and may be associated with hardware such as a GPU installed in one or more servers **312**, whereas firmware **602-3** may correspond to firmware Z and may be associated with High Density memory (HDM) installed in one or more servers **312**. Each firmware **602** may be mapped one-to-one with a proxy pod **606** in a proxy node **604** and a redundant proxy pod **608** in a redundant proxy node **610** (e.g., firmware **602-1** may be mapped to proxy pod **606-1** in proxy node **604** and redundant proxy pod **608-1** in redundant proxy node **610**).

[0077] Redundant proxy node **610** may make sure connectivity is maintained with server **312**. For example, proxy node **604** may be on a first server **312** and redundant proxy node **610** may be on a second server **312**. Thus, if firmware **602** needs to be updated or a server **312** needs to be powered on and proxy node **604** is down, firmware **602** may be updated using redundant proxy node **610**.

[0078] Proxy node **604** and redundant proxy node **610** may form part of proxy cluster **612**. Proxy clusters **612** may be configured to manage servers **312** hosting proxy nodes **604** and redundant proxy nodes **610**. For example, a proxy cluster **612** may be configured such that a server **312** hosting proxy node **604** and a server **312** hosting redundant proxy node **610** are always powered on, a server **312** hosting proxy node **604** is powered on when a server **312** hosting redundant proxy node **610** is powered down, idle or not responding, or a server **312** hosting redundant proxy node **610** is powered on when a server **312** hosting proxy node **604** is powered down, idle or not responding.

[0079] Embodiments may also provide redundancy in available services **350**. Cloud services server **216** may contain service cluster **622** with service pods **616-1** to **616-S** in service node **614** and service pods **618-1** to **618-S** in redundant service node **620**. Presently, there are approximately **200** available services **350** that may be provisioned using embodiments disclosed herein. Thus, if an available service **350** cannot be provisioned using a service pod **616** (e.g., a server **312** hosting service node **614** is powered down or idle), the available service **350** may still be provisioned



using redundant service pod **618**. For example, service pod **616-1** may correspond with credentials service **350-15** (depicted in FIG. 5). If a firmware update is needed to provision credentials service **350-15** to one or more servers **312**, embodiments may try to provision credentials service **350-15** using a first firmware update service (e.g., service 1 pod **616-1**). If service pod **616-1** is unavailable, embodiments may provision credentials service **350-15** using redundant service 1 pod **618-1**.

[0080] Referring back to FIG. 3, access cluster **624** may contain group nodes **626-1** to **626-M**, wherein all servers **312** in a group **320** are associated with a group node **626**. For example, servers **312** in group **320-1** depicted in FIG. 3 are associated with group 1 node **626-1** in FIG. 6. If an available service **350** is to be provisioned to servers **312** in group **320-1**, embodiments may provision the available service **350** using group 1 node **626-1**. A group node **626** may be associated with one or more servers **312** based on hardware or node access. Servers **312** can be added to or removed from group nodes **626** as needed.

#### Telemetry Streaming

[0081] Traditionally, each RAC **14** facilitates registering the metric resource definitions (MRD). In some servers **12** and data centers **10**, there can be overlapping duplicate definitions where two or more RACs **14** stream redundant telemetry data over a network. Furthermore, for any telemetry event, each server **12** must register for an individual event. If multiple servers **12** register for the same event, RAC **14** often sends the telemetry data, resulting in duplicate telemetry data. As a result, telemetry data may not be sent or may be incomplete. For example, many RACs **14** have limited available network and storage bandwidth. Often, if a RAC **14** cannot push out the event, then RAC **14** discards the event, even if the event is a critical event. If many end user information handling systems register for different telemetry events to a single RAC **14**, the RAC **14** may try to send the same telemetry data to multiple end user information handling systems. Even due to user preferences, RAC network bandwidth may also be used to send duplicate telemetry data to multiple end user information handling systems. At some point, RAC network and storage hardware will reach a maximum usable limit. After that, to support more traffic and for scalability, new hardware must be added to servers **12** to scale for given traffic and data, which will ultimately cost users. After some point of time this scaling also will be overloaded and after that, data centers **10** and servers **12** will not be able to add any new hardware to support further scaling.

#### Baseboard Management Controllers do not Adjust Telemetry Streaming

[0082] Currently, BMCs in servers **12** may have different streaming rates for different metrics. Furthermore, BMCs in servers **12** lack intelligence to adjust telemetry collection based on server events/errors. Proactive collection of server health state may be advantageous for analytics and diagnostics.

[0083] Many RACs **14** do not support open telemetry (OTEL). To support OTEL, firmware in RAC **14** needs to be released for new features **30**. Similarly, if any new protocol/schema is released, firmware in RAC **14** needs to be released.

#### Common Scenarios for Data Loss in Telemetry

##### Hardware Interconnect Access Threshold

[0084] Data collection may be limited. For example, Platform Environmental Control Interface-based data collection is generally limited by I<sup>2</sup>C bandwidth (e.g., 100 Khz-400 Khz) as well as the number of transactions occurring at an I<sup>2</sup>C bus for Control Plane and User Plane Separation (CUPS) data (e.g., 300 commands). A data collection requirement may indicate CPU data is to be collected every 5 seconds. However, due to multiple sources using the I<sup>2</sup>C bus to access the CPU data, the CPU data may actually be collected only every 8 seconds, resulting in a loss of telemetry data such that the data collection requirement is not met.

##### Missing Data Due to Network Glitch or Overload

[0085] A set of telemetry data might have missing data if there are frequent packet drops over a network. As a result, a customer information handling system communicatively coupled to servers **12** in data center **10** and running any learning algorithm may need to address missing telemetry data with an algorithm, which adds overhead. Also, a telemetry data mismatch may occur between different information handling systems communicatively coupled to servers **12** in data center **10** since they independently collect telemetry data.

##### Embodiments Analyze Telemetry Data to Reduce Network and Storage Bandwidth but Still Meet Collection Requirements

[0086] Embodiments disclosed herein may comprise a system configured such that a BMC or RAC **314** sends telemetry data to private cloud server **60** instead of external cloud servers or customer systems, which frees up bandwidth and storage. In some embodiments, RAC **314** in each server **312** may send telemetry data to PCS **360** but not to external cloud server **210**. Services **350** executing in PCS **360** may receive all the telemetry data sent by one or more server **312**. Advantageously, PCS **360** knows the data collection requirements for each customer information handling system. As a result, PCS **360** may receive and record a first set of telemetry data for one or more RACs **314** and communicate a second set of telemetry data to one or more customer information handling systems, wherein the second set of telemetry data sent to the one or more customer information handling systems has less data than the first set of telemetry data sent to PCS **360**. For example, a first customer information handling system may require collection of telemetry data corresponding to a CPU and require the telemetry data be sent every 5 seconds and a second customer information handling system may require the same telemetry data but require the telemetry data be sent every 8 seconds. RAC **314** may send telemetry data every 5 seconds and every 8 seconds, but PCS **360** may be configured to communicate only the telemetry data sent every 5 seconds because that set of telemetry data meets the collection requirements for both customer information handling systems and sending only that set of telemetry data may reduce network and storage bandwidth.

[0087] Referring to FIG. 7, embodiments of a system for telemetry streaming may include RAC **314** comprising firmware storing Hardware Abstraction Layer (HAL) **22**



including power control **20-1** and thermal control **20-2**, as well as Inter-Integrated Circuit (I2C) protocol Library (LIBI2C) firmware **22-1** and management component transport protocol (MCTP) firmware **22-6**, as described above. Each RAC **314** may further comprise firmware such as D-Bus interface **18**, Redfish firmware **30-2** and Telemetry firmware **30-15**, PRC **302** (which may be a google Remote Procedure Call queue), and D-Bus to PRC mapper **304**, as described above.

[0088] Referring to FIG. 4 and FIG. 7, in addition to available services **350** such as cloud services API **350-1**, Redfish service **350-2**, hardware inventory service **350-5**, NC-SI service **350-10**; PLDM service **350-11**; SPDML service **350-12** described relative to FIG. 4, private cloud server **360** may comprise firmware storing other available services **350**.

[0089] Telemetry streaming service **350-27** may refer to firmware for communicating telemetry data to external cloud servers **210**.

[0090] Redfish-Server Sent Event (SSE) service **350-28** may refer to an available service **350** that continuously pushes rapidly changing telemetry data.

[0091] Telemetry configuration service **350-29** may monitor requests for telemetry collection and responses to the requests to determine a telemetry configuration for private cloud server **360**. For example, if telemetry configuration service **350-29** determines multiple customer information handling systems are sending requests for telemetry data corresponding to CPU data, telemetry configuration service **350-31** may determine a minimum collection requirement for CPU data and configure PCS **360** to communicate telemetry data that meets or exceeds the minimum collection requirement.

[0092] Error telemetry configuration service **350-31** may communicate with one or more available services **350** to identify errors in telemetry collection and communication. In some embodiments, error telemetry configuration service **350-31** may receive multiple instances of telemetry error and determine the multiple instances represent the same telemetry error.

[0093] Telemetry collection analytics service **350-32** may analyze what telemetry data is being collected to identify required telemetry data, error telemetry data and optional telemetry data.

[0094] Attribute registry **350-33** may store information that defines a property of each file corresponding to telemetry data.

[0095] Aggregator service **350-34** may determine what telemetry data can be aggregated and store the aggregated telemetry data in centralized telemetry database **350-30** or localized telemetry database **332**. For example, if multiple RACs **314** are sending telemetry data and a data center experiences an issue with a power source or a network connectivity issue, each RAC **314** may send error data. Instead of each RAC **314** sending multiple instances of error data to multiple customer information handling systems, aggregator service **350-34** may determine all the error data relates to the same error source and send a single error flag to all the customer information handling systems, reducing network bandwidth.

[0096] Filter service **350-37** may determine what telemetry data can be filtered and store the filtered telemetry data in centralized telemetry database **350-30** or localized telem-

etry database **332**. For example, filter service **350-37** may remove duplicate telemetry data.

[0097] Jaeger service **350-38** may refer to software used to monitor and troubleshoot problems on interconnected software components (e.g., available services **350**) that communicate with each other to execute a software function.

[0098] Some embodiments may use one or more external cloud servers **210** as an intermediary service provider, wherein client information handling systems communicatively coupled to external cloud servers **210** may request telemetry data from external cloud server **210** and not directly to RAC **314** or server **312**. In this scenario, RAC **314** would not need to know any telemetry collection requirements, (e.g., what telemetry data is required by a client information handling system, how many information handling systems are requesting the same telemetry data, etc.).

[0099] Advantageously, RAC **314** may continuously send telemetry data irrespective of a telemetry request to private cloud server **360** and private cloud server **360** may communicate telemetry data to external cloud server **210**, which has higher storage, network and processing power, and can directly interact with users and respond to requests.

[0100] FIG. 8 is a flow diagram depicting an embodiment of a method **800** for streaming telemetry data.

[0101] At step **802**, embodiments may begin collecting telemetry data for a server **312**. The telemetry data may be based on a default configuration.

[0102] Embodiments may define three types of config files which enables the telemetry collections. These types may be, for example, a mandatory list, an optional list and an error list. Mandatory list config files may be based on the understanding the telemetry collection requirements of a customer. Mandatory list config files may be based on any learning mechanism such as time series or frequent occurring patterns. Optional list config files may comprise any remaining telemetry data. Error list config files may comprise telemetry data groups for different types of errors which are only run when one of error type occurs. This enables better debugging post error. The lists are subject to change and may be adapted to meet a user collection requirement and may be adjustable to frequency of the collection.

[0103] In some embodiments, new servers **312** may be configured to collect telemetry data based on existing servers' **312** configurations and reconfigured based on new telemetry collection requirements.

[0104] RAC **314** may continuously collect telemetry data, while pushing the data to centralized telemetry database **330** based on a RAC service-defined policy. In case there is a critical event, localized telemetry database **332** may provide useful telemetry data that was not required by the policy, which may be advantageous for predictive collection of important telemetry data.

[0105] At step **804**, embodiments may monitor telemetry data collection. Embodiments may have telemetry logic for understanding the outgoing telemetry traffic by monitoring the type of queries RAC **314** is servicing related to telemetry measurements. The adaptive telemetry collection may have telemetry measurements/data enabled that are related to, for example, monitoring of errors and failures, based on Redfish/Web-Service Management (WSMAN) queries and telemetry data related to ports.



[0106] At step 806, embodiments may determine if there is more telemetry data to be collected. Determining if more telemetry data is to be collected may include determining a new information handling system is requesting telemetry data or the frequency at which telemetry data is to be collected has increased.

[0107] At step 808, embodiments may determine if a new configuration is required to meet one or more telemetry collection requirements. A new configuration may include, for example, that more telemetry data is to be collected, that telemetry data should be sent at a higher frequency, and/or is to be sent to more customer information handling systems. Some embodiments may enhance or reduce the frequency at which telemetry data is sent to private cloud server 360 based on a warning/error on a server 312 or based on the data center admin's action/settings. For example, if data center 310 is experiencing network issues, embodiments may store telemetry data in centralized telemetry database and decrease the frequency at which telemetry data is sent to external cloud servers 210. Once the network issues are resolved, embodiments may increase the frequency at which telemetry data is sent to external cloud servers 210.

[0108] At step 810, if a new configuration is needed, embodiments may check one or more telemetry data groups to see what telemetry data is being sent. In some embodiments, telemetry data may be filtered before being communicated to external cloud servers 210.

[0109] A step 812, if a new configuration is needed, embodiments may check the frequency at which telemetry data is being sent.

[0110] A step 814, if a new configuration is needed, embodiments may check if telemetry data can be consolidated or aggregated.

[0111] Based on steps 810-814, embodiments may collect, filter, aggregate and store telemetry data received from multiple servers 312 in centralized telemetry database 330 or localized telemetry database 332 before being communicated to external cloud servers 210. A set of telemetry data communicated to one or more external cloud servers 210 may include less data than a set of telemetry data received by private cloud server 360 but still meet telemetry data collection requirements.

[0112] By aggregating and filtering telemetry data and communicating a single set of telemetry data to external cloud server 210 only and not to end user information handling systems, network bandwidth may be used to process a single instance of telemetry data, minimizing RAC bandwidth. Furthermore, private cloud server 360 can store telemetry data such that RAC 314 can save storage for its own data and client related data can be stored in external cloud server 210.

[0113] Advantageously, embodiments may increase a service life of RAC 314 and save customer money to not to spend more on adding new hardware, instead get more from external cloud server 210 and over all that will lead to benefit to end users.

[0114] In some embodiments, telemetry streaming service 350-27 may collect all telemetry data and communicate the telemetry data to external cloud server 210, wherein external cloud servers 210 may store the telemetry data and multiple telemetry services on external cloud servers 210 can read the telemetry data from external cloud servers 210 and process the telemetry data at a higher rate. For example, Prometheus service 720-1 is an example of a systems and service

monitoring system that may be stored on external cloud server 210. Prometheus service 720-1 may collect metrics from external cloud server 210 related to configured targets at given intervals, evaluate rule expressions, display the results, and can trigger alerts when specified conditions are observed. Grafana service 720-2 is an example of an open-source analytics and monitoring service.

[0115] The above disclosed subject matter is to be considered illustrative, and not restrictive, and the appended claims are intended to cover all such modifications, enhancements, and other embodiments which fall within the true spirit and scope of the present disclosure. Thus, to the maximum extent allowed by law, the scope of the present disclosure is to be determined by the broadest permissible interpretation of the following claims and their equivalents, and shall not be restricted or limited by the foregoing detailed description.

What is claimed is:

1. A system for telemetry streaming, the system comprising:
  - a remote access controller (RAC) on a server, the RAC comprising:
    - a RAC processor; and
    - a RAC memory storing a set of instructions executable by the RAC processor to collect a first set of telemetry data associated with a plurality of devices on the server based on one or more telemetry collection requirements;
  - a private cloud server (PCS) comprising:
    - a PCS processor; and
    - a PCS memory storing a set of available services and a set of instructions executable by the PCS processor to:
      - communicate with the RAC to receive the first set of telemetry data from the RAC;
      - communicate with the RAC to monitor outgoing telemetry traffic associated with the first set of telemetry data;
      - store the first set of telemetry data in a telemetry database;
      - determine the one or more telemetry collection requirements; and
      - communicate a second set of telemetry data to an external cloud server, wherein the second set of telemetry data communicated to the external cloud server contains less data than the first set of telemetry data and meets the one or more telemetry collection requirements.
2. The system of claim 1, wherein the PCS memory stores a telemetry filter service configured to filter the first set of telemetry data.
3. The system of claim 1, wherein the PCS memory stores a telemetry aggregator service configured to aggregate the first set of telemetry data into the second set of telemetry data.
4. The system of claim 1, wherein the PCS memory stores a telemetry streaming service configured to determine a frequency of communicating the second set of telemetry data to the external cloud server.
5. The system of claim 1, wherein the set of PCS instructions are executable by the PCS processor for communicating the second set of telemetry data to the external cloud server based on the second set of telemetry data corresponding to a mandatory list of telemetry data.



6. The system of claim 1, wherein the set of PCS instructions are executable by the PCS processor for communicating the second set of telemetry data to the external cloud server based on the second set of telemetry data corresponding to an error list of telemetry data.

7. The system of claim 1, wherein the set of PCS instructions are executable by the PCS processor for communicating the second set of telemetry data to the external cloud server based on the second set of telemetry data corresponding to an optional list of telemetry data.

8. A data center comprising:

a plurality of servers, wherein each server comprises:

a remote access controller (RAC) comprising:

a RAC processor; and

a RAC memory storing a set of instructions executable by the RAC processor to collect a first set of telemetry data associated with the plurality of devices on the server based on one or more telemetry collection requirements;

wherein at least one server of the plurality of servers comprises a private cloud server (PCS) communicatively coupled to the RAC in each server of the plurality of servers and an external cloud server, the private cloud server comprising:

a PCS processor; and

a PCS memory storing a set of available services and a set of instructions executable by the PCS processor to:

communicate with the RAC to receive the first set of telemetry data from the RAC in each server of the plurality of servers;

communicate with the RAC in each server of the plurality of servers to monitor outgoing telemetry traffic associated with the first set of telemetry data;

store the first set of telemetry data in a telemetry database;

determine the one or more telemetry collection requirements; and

communicate a second set of telemetry data to the external cloud server, wherein the second set of telemetry data communicated to the external cloud server contains less data than the first set of telemetry data and meets the one or more telemetry collection requirements.

9. The data center of claim 8, wherein the PCS memory stores a telemetry filter service configured to filter the first set of telemetry data.

10. The data center of claim 8, wherein the PCS memory stores a telemetry aggregator service for aggregating the first set of telemetry data received from the plurality of servers into the second set of telemetry data.

11. The data center of claim 8, wherein the PCS memory stores a telemetry streaming service for determining a frequency of communicating the second set of telemetry data to the external cloud server.

12. The data center of claim 8, wherein the set of PCS instructions are executable by the PCS processor for com-

municating the second set of telemetry data to the external cloud server based on the second set of telemetry data corresponding to a mandatory list of telemetry data.

13. The data center of claim 8, wherein the set of PCS instructions are executable by the PCS processor for communicating the second set of telemetry data to the external cloud server based on the second set of telemetry data corresponding to an error list of telemetry data.

14. The data center of claim 8, wherein the set of PCS instructions are executable by the PCS processor for communicating the second set of telemetry data to the external cloud server based on the second set of telemetry data corresponding to an optional list of telemetry data.

15. A method of operating a data center comprising a plurality of servers, the method comprising:

storing, on each server of the plurality of servers, a set of instructions executable by the RAC processor to collect a first set of telemetry data associated with a plurality of devices on the server based on one or more telemetry collection requirements;

executing, by a private cloud server communicatively coupled to the RAC in each server of the plurality of servers, a set of instructions to perform:

communicating with the RAC in each server of the plurality of servers to receive the first set of telemetry data from the RAC in each server of the plurality of servers;

communicating with the RAC in each server of the plurality of servers to monitor outgoing telemetry traffic associated with the first set of telemetry data;

storing the first set of telemetry data received from the plurality of servers in a telemetry database;

determining one or more telemetry collection requirements; and

communicating a second set of telemetry data to an external cloud server, wherein the second set of telemetry data communicated to the external cloud server contains less data than the first set of telemetry data and meets the one or more telemetry collection requirements.

16. The method of claim 15, further comprising filtering the first set of telemetry data received from the plurality of servers into the second set of telemetry data.

17. The method of claim 15, further comprising aggregating the first set of telemetry data received from the plurality of servers into the second set of telemetry data.

18. The method of claim 15, further comprising determining a frequency of communicating the second set of telemetry data to the external cloud server.

19. The method of claim 15, further comprising communicating the second set of telemetry data to the external cloud server based on the second set of telemetry data corresponding to one of a mandatory list of telemetry data, an error list of telemetry data or an optional list of telemetry data.

\* \* \* \* \*